
What Makes Data Suitable for a Locally Connected Neural Network? A Necessary and Sufficient Condition Based on Quantum Entanglement

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 The question of what makes a data distribution suitable for deep learning is a funda-
2 mental open problem. Focusing on locally connected neural networks (a prevalent
3 family of architectures that includes convolutional and recurrent neural networks as
4 well as local self-attention models), we address this problem by adopting theoretical
5 tools from quantum physics. Our main theoretical result states that a certain locally
6 connected neural network is capable of accurate prediction over a data distribution
7 *if and only if* the data distribution admits low quantum entanglement under certain
8 canonical partitions of features. As a practical application of this result, we derive a
9 preprocessing method for enhancing the suitability of a data distribution to locally
10 connected neural networks. Experiments with widespread models over various
11 datasets demonstrate our findings. We hope that our use of quantum entanglement
12 will encourage further adoption of tools from physics for formally reasoning about
13 the relation between deep learning and real-world data.

14 1 Introduction

15 Deep learning is delivering unprecedented performance when applied to data modalities involving
16 images, text and audio. On the other hand, it is known both theoretically and empirically [52, 1] that
17 there exist data distributions over which deep learning utterly fails. The question of *what makes a*
18 *data distribution suitable for deep learning* is a fundamental open problem in the field.

19 A prevalent family of deep learning architectures is that of *locally connected neural networks*. It
20 includes, among others: (i) convolutional neural networks, which dominate the area of computer
21 vision; (ii) recurrent neural networks, which were the most common architecture for sequence
22 (e.g. text and audio) processing, and are experiencing a resurgence by virtue of S4 models [26]; and
23 (iii) local variants of self-attention neural networks [46]. Conventional wisdom postulates that data
24 distributions suitable for locally connected neural networks are those exhibiting a “local nature,”
25 and there have been attempts to formalize this intuition [64, 28, 14]. However, to the best of our
26 knowledge, there are no characterizations providing necessary and sufficient conditions for a data
27 distribution to be suitable to a locally connected neural network.

28 A seemingly distinct scientific discipline tying distributions and computational models is *quantum*
29 *physics*. There, distributions of interest are described by *tensors*, and the associated computational
30 models are *tensor networks*. While there is shortage in formal tools for assessing the suitability of
31 data distributions to deep learning architectures, there exists a widely accepted theory that allows for
32 assessing the suitability of tensors to tensor networks. The theory is based on the notion of *quantum*
33 *entanglement*, which quantifies dependencies that a tensor admits under partitions of its axes (for
34 a given tensor \mathcal{A} and a partition of its axes to sets \mathcal{K} and \mathcal{K}^c , the entanglement is a non-negative
35 number quantifying the dependence that \mathcal{A} induces between \mathcal{K} and \mathcal{K}^c).

36 In this paper, we apply the foregoing theory to a tensor network equivalent to a certain locally
 37 connected neural network, and derive theorems by which fitting a tensor is possible if and only
 38 if the tensor admits low entanglement under certain *canonical partitions* of its axes. We then
 39 consider the tensor network in a machine learning context, and find that its ability to attain low
 40 approximation error, *i.e.* to express a solution with low population loss, is determined by its ability to
 41 fit a particular tensor defined by the data distribution, whose axes correspond to features. Combining
 42 the latter finding with the former theorems, we conclude that a *locally connected neural network is*
 43 *capable of accurate prediction over a data distribution if and only if the data distribution admits low*
 44 *entanglement under canonical partitions of features*. Experiments with different datasets corroborate
 45 this conclusion, showing that the accuracy of common locally connected neural networks (including
 46 modern convolutional, recurrent, and local self-attention neural networks) is inversely correlated to
 47 the entanglement under canonical partitions of features in the data (the lower the entanglement, the
 48 higher the accuracy, and vice versa).

49 The above results bring forth a recipe for enhancing the suitability of a data distribution to locally
 50 connected neural networks: given a dataset, search for an arrangement of features which leads to low
 51 entanglement under canonical partitions, and then arrange the features accordingly. Unfortunately,
 52 the above search is computationally prohibitive. However, if we employ a certain correlation-
 53 based measure as a surrogate for entanglement, *i.e.* as a gauge for dependence between sides of a
 54 partition of features, then the search converts into a succession of *minimum balanced cut* problems,
 55 thereby admitting use of well-established graph theoretical tools, including ones designed for large
 56 scale [29, 56]. We empirically evaluate this approach on various datasets, demonstrating that it
 57 substantially improves prediction accuracy of common locally connected neural networks (including
 58 modern convolutional, recurrent, and local self-attention neural networks).

59 The data modalities to which deep learning is most commonly applied — namely ones involving
 60 images, text and audio — are often regarded as natural (as opposed to, for example, tabular data
 61 fusing heterogeneous information). We believe the difficulty in explaining the suitability of such
 62 modalities to deep learning may be due to a shortage in tools for formally reasoning about natural
 63 data. Concepts and tools from physics — a branch of science concerned with formally reasoning
 64 about natural phenomena — may be key to overcoming said difficulty. We hope that our use of
 65 quantum entanglement will encourage further research along this line.

66 2 Preliminaries

67 For simplicity, the main text treats locally connected neural networks whose input data is one
 68 dimensional (*e.g.* text and audio). We defer to Appendix I an extension of the analysis and experiments
 69 to models intaking data of arbitrary dimension (*e.g.* two-dimensional images). Due to lack of space,
 70 we also defer our review of related work to Appendix A.

71 We use $\|\cdot\|$ and $\langle \cdot, \cdot \rangle$ to denote the Euclidean (Frobenius) norm and inner product, respectively.
 72 We shorthand $[N] := \{1, \dots, N\}$, where $N \in \mathbb{N}$. The complement of $\mathcal{K} \subseteq [N]$ is denoted by
 73 $\mathcal{K}^c := [N] \setminus \mathcal{K}$.

74 2.1 Tensors and Tensor Networks

75 For our purposes, a *tensor* is a multi-dimensional array $\mathcal{A} \in \mathbb{R}^{D_1 \times \dots \times D_N}$, where $N \in \mathbb{N}$ is its
 76 *dimension* and $D_1, \dots, D_N \in \mathbb{N}$ are its *axes lengths*. The (d_1, \dots, d_N) 'th entry of \mathcal{A} is denoted
 77 $\mathcal{A}_{d_1, \dots, d_N}$.

78 *Contraction* between tensors is a generalization of multiplication between matrices. Two matrices
 79 $\mathbf{A} \in \mathbb{R}^{D_1 \times D_2}$ and $\mathbf{B} \in \mathbb{R}^{D'_1 \times D'_2}$ can be multiplied if $D_2 = D'_1$, in which case we get a matrix in
 80 $\mathbb{R}^{D_1 \times D'_2}$ holding $\sum_{d=1}^{D_2} \mathbf{A}_{d_1, d} \cdot \mathbf{B}_{d, d'_2}$ in entry $(d_1, d'_2) \in [D_1] \times [D'_2]$. More generally, two tensors
 81 $\mathcal{A} \in \mathbb{R}^{D_1 \times \dots \times D_N}$ and $\mathcal{B} \in \mathbb{R}^{D'_1 \times \dots \times D'_{N'}}$ can be contracted along axis $n \in [N]$ of \mathcal{A} and $n' \in [N']$
 82 of \mathcal{B} if $D_n = D'_{n'}$, in which case we get a tensor of size $D_1 \times \dots \times D_{n-1} \times D_{n+1} \times \dots \times D_N \times D'_1 \times$
 83 $\dots \times D'_{n'-1} \times D'_{n'+1} \times \dots \times D'_{N'}$ holding $\sum_{d=1}^{D_n} \mathcal{A}_{d_1, \dots, d_{n-1}, d, d_{n+1}, \dots, d_N} \cdot \mathcal{B}_{d'_1, \dots, d'_{n'-1}, d, d'_{n'+1}, \dots, d'_{N'}}$
 84 in the entry indexed by $\{d_k \in [D_k]\}_{k \in [N] \setminus \{n\}}$ and $\{d'_k \in [D'_k]\}_{k \in [N'] \setminus \{n'\}}$.

85 *Tensor networks* are prominent computational models for fitting (*i.e.* representing) tensors. More
 86 specifically, a tensor network is a weighted graph that describes formation of a (typically high-
 87 dimensional) tensor via contractions between (typically low-dimensional) tensors. As customary

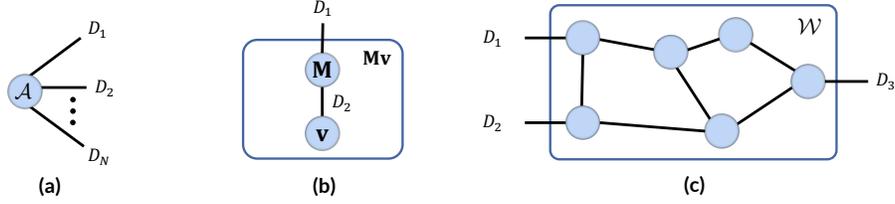


Figure 1: Tensor networks form a graphical language for fitting (*i.e.* representing) tensors through tensor contractions. **Tensor network definition:** Every node in a tensor network is associated with a tensor, whose dimension is equal to the number of edges emanating from the node. An edge connecting two nodes specifies contraction between the tensors associated with the nodes (Section 2.1), where the weight of the edge signifies the respective axes lengths. Tensor networks may also contain open edges, *i.e.* edges that are connected to a node on one side and are open on the other. The number of such open edges is equal to the dimension of the tensor produced by contracting the tensor network. **Illustrations:** Presented are exemplar tensor network diagrams of: (a) an N -dimensional tensor $\mathcal{A} \in \mathbb{R}^{D_1 \times \dots \times D_N}$; (b) a vector-matrix multiplication between $\mathbf{M} \in \mathbb{R}^{D_1 \times D_2}$ and $\mathbf{v} \in \mathbb{R}^{D_2}$, which results in the vector $\mathbf{M}\mathbf{v} \in \mathbb{R}^{D_1}$; and (c) a tensor network generating $\mathcal{W} \in \mathbb{R}^{D_1 \times D_2 \times D_3}$.

88 (*cf.* [42]), we will present tensor networks via graphical diagrams to avoid cumbersome notation —
 89 see Figure 1 for details.

90 2.2 Quantum Entanglement

91 In quantum physics, the distribution of possible states for a multi-particle (“many body”) system is
 92 described by a tensor, whose axes are associated with individual particles. A key property of the
 93 distribution is the dependence it admits under a given partition of the particles (*i.e.* between a given
 94 set of particles and its complement). This dependence is formalized through the notion of *quantum*
 95 *entanglement*, defined using the distribution’s description as a tensor — see Definition 1 below.

96 Quantum entanglement lies at the heart of a widely accepted theory which allows assessing the ability
 97 of a tensor network to fit a given tensor (*cf.* [15, 35]). In Section 3 we specialize this theory to a
 98 tensor network equivalent to a certain locally connected neural network.

99 **Definition 1.** For a tensor $\mathcal{A} \in \mathbb{R}^{D_1 \times \dots \times D_N}$ and subset of its axes $\mathcal{K} \subseteq [N]$, let $\llbracket \mathcal{A}; \mathcal{K} \rrbracket \in$
 100 $\mathbb{R}^{\prod_{n \in \mathcal{K}} D_n \times \prod_{n \in \mathcal{K}^c} D_n}$ be the arrangement of \mathcal{A} as a matrix where rows correspond to axes \mathcal{K} and
 101 columns correspond to the remaining axes $\mathcal{K}^c := [N] \setminus \mathcal{K}$. Denote by $\sigma_1 \geq \dots \geq \sigma_{D_{\mathcal{K}}} \in \mathbb{R}_{\geq 0}$ the sin-
 102 gular values of $\llbracket \mathcal{A}; \mathcal{K} \rrbracket$, where $D_{\mathcal{K}} := \min\{\prod_{n \in \mathcal{K}} D_n, \prod_{n \in \mathcal{K}^c} D_n\}$. The *quantum entanglement*¹ of
 103 \mathcal{A} with respect to the partition $(\mathcal{K}, \mathcal{K}^c)$ is the entropy of the distribution $\{\rho_d := \sigma_d^2 / \sum_{d'=1}^{D_{\mathcal{K}}} \sigma_{d'}^2\}_{d=1}^{D_{\mathcal{K}}}$,
 104 *i.e.* $QE(\mathcal{A}; \mathcal{K}) := -\sum_{d=1}^{D_{\mathcal{K}}} \rho_d \ln(\rho_d)$. By convention, if $\mathcal{A} = 0$ then $QE(\mathcal{A}; \mathcal{K}) = 0$.

105 3 Low Entanglement Under Canonical Partitions Is Necessary and 106 Sufficient for Fitting Tensor

107 In this section, we prove that a tensor network equivalent to a certain locally connected neural network
 108 can fit a tensor if and only if the tensor admits low entanglement under certain canonical partitions of
 109 its axes. We begin by introducing the tensor network (Section 3.1). Subsequently, we establish the
 110 necessary and sufficient condition required for it to fit a given tensor (Section 3.2). For conciseness,
 111 the treatment in this section is limited to one-dimensional (sequential) models; see Appendix I.1 for
 112 an extension to arbitrary dimensions.

113 3.1 Tensor Network Equivalent to a Locally Connected Neural Network

114 Let $N \in \mathbb{N}$, and for simplicity suppose that $N = 2^L$ for some $L \in \mathbb{N}$. We consider a tensor network
 115 with an underlying perfect binary tree graph of height L , which generates $\mathcal{W}_{\text{TN}} \in \mathbb{R}^{D_1 \times \dots \times D_N}$.
 116 Figure 2(a) provides its diagrammatic definition. For simplicity, the lengths of axes corresponding to
 117 inner (non-open) edges are taken to be $R \in \mathbb{N}$, referred to as the *width* of the tensor network.

118 As identified by previous works, the tensor network depicted in Figure 2(a) is equivalent to a certain
 119 locally connected neural network (with polynomial non-linearity — see, *e.g.*, [11, 9, 35, 49]). In
 120 particular, contracting the tensor network with vectors $\mathbf{x}^{(1)} \in \mathbb{R}^{D_1}, \dots, \mathbf{x}^{(N)} \in \mathbb{R}^{D_N}$, as illustrated
 121 in Figure 2(b), can be viewed as a forward pass of the data instance $(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)})$ through a locally

¹There exist multiple notions of entanglement in quantum physics (see, *e.g.*, [35]). The one we consider is the most common, known as *entanglement entropy*.

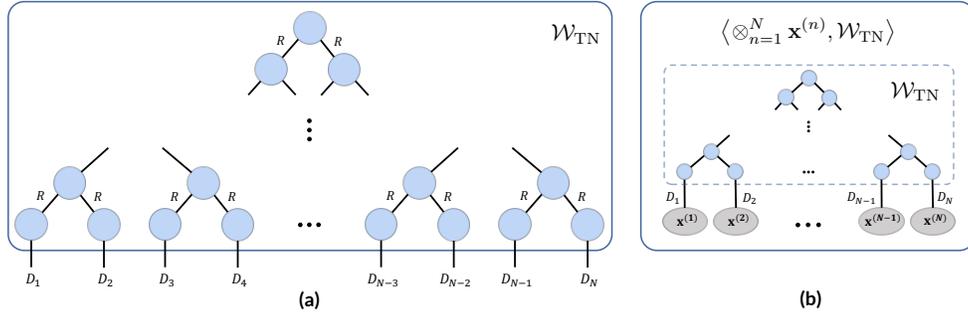


Figure 2: The analyzed tensor network equivalent to a locally connected neural network. **(a)** We consider a tensor network adhering to a perfect binary tree connectivity with $N = 2^L$ leaf nodes, for $L \in \mathbb{N}$, generating $\mathcal{W}_{\text{TN}} \in \mathbb{R}^{D_1 \times \dots \times D_N}$. Axes corresponding to open edges are indexed such that open edges descendant to any node of the tree have contiguous indices. The lengths of axes corresponding to inner (non-open) edges are equal to $R \in \mathbb{N}$, referred to as the *width* of the tensor network. **(b)** Contracting \mathcal{W}_{TN} with vectors $\mathbf{x}^{(1)} \in \mathbb{R}^{D_1}, \dots, \mathbf{x}^{(N)} \in \mathbb{R}^{D_N}$ produces $\langle \otimes_{n=1}^N \mathbf{x}^{(n)}, \mathcal{W}_{\text{TN}} \rangle$. Performing these contractions from leaves to root can be viewed as a forward pass of a data instance $(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)})$ through a certain locally connected neural network (with polynomial non-linearity; see, e.g., [11, 9, 35, 49]). Accordingly, we call the tensor network generating \mathcal{W}_{TN} a *locally connected tensor network*.

122 connected neural network. This computation results in a scalar equal to $\langle \otimes_{n=1}^N \mathbf{x}^{(n)}, \mathcal{W}_{\text{TN}} \rangle$, where \otimes
 123 stands for the outer product.² In light of its equivalence to a locally connected neural network, we
 124 will refer to the tensor network as a *locally connected tensor network*. We note that for the equivalent
 125 neural network to be practical (in terms of memory and runtime), the width of the tensor network R
 126 needs to be of moderate size. Specifically, R cannot be exponential in the dimension N , meaning
 127 $\ln(R)$ needs to be much smaller than N .

128 By virtue of the locally connected tensor network’s equivalence to a deep neural network, it has been
 129 paramount for the study of expressiveness and generalization in deep learning [11, 8, 9, 12, 13, 53, 34,
 130 35, 2, 30, 31, 36, 47, 48, 49, 50]. Although the equivalent deep neural network (which has polynomial
 131 non-linearity) is less common than other neural networks (e.g., ones with ReLU non-linearity), it
 132 has demonstrated competitive performance in practice [7, 10, 54, 57, 22]. More importantly, its
 133 analyses (through its equivalence to the locally connected tensor network) brought forth numerous
 134 insights that were demonstrated empirically and led to development of practical tools for common
 135 locally connected architectures. Continuing this line, we will demonstrate our theoretical insights
 136 through experiments with widespread convolutional, recurrent and local self-attention architectures
 137 (Section 4.3), and employ our theory for deriving an algorithm that enhances the suitability of a data
 138 distribution to said architectures (Section 5).

139 3.2 Necessary and Sufficient Condition for Fitting Tensor

140 Herein we show that the ability of the locally connected tensor network (defined in Section 3.1) to fit
 141 (i.e. represent) a given tensor is determined by the entanglements that the tensor admits under the
 142 following *canonical partitions* of $[N]$.

143 **Definition 2.** The *canonical partitions* of $[N]$ (illustrated in Figure 4 of Appendix B) are:

$$\mathcal{C}_N := \left\{ (\mathcal{K}, \mathcal{K}^c) : \mathcal{K} = \{2^{L-l} \cdot (n-1) + 1, \dots, 2^{L-l} \cdot n\}, l \in \{0, \dots, L\}, n \in [2^l] \right\}.$$

144 By appealing to known upper bounds on the entanglements that a given tensor network supports [15,
 145 35], we establish that if the locally connected tensor network can fit a given tensor, that tensor must
 146 admit low entanglement under the canonical partitions of its axes. Namely, suppose that \mathcal{W}_{TN} — the
 147 tensor generated by the locally connected tensor network — well-approximates an N -dimensional
 148 tensor \mathcal{A} . Then, Theorem 1 below shows that the entanglement of \mathcal{A} with respect to a canonical
 149 partition cannot be much larger than $\ln(R)$ (recall that R is the width of the locally connected tensor
 150 network), whereas the entanglement attainable by an arbitrary tensor with respect to a canonical
 151 partition can be linear in the dimension N .

152 In the other direction, Theorem 2 below implies that low entanglement under the canonical partitions
 153 is not only necessary for a tensor to be fit by the locally connected tensor network, but also sufficient.

²For any $\{\mathbf{x}^{(n)} \in \mathbb{R}^{D_n}\}_{n=1}^N$, the outer product $\otimes_{n=1}^N \mathbf{x}^{(n)} \in \mathbb{R}^{D_1 \times \dots \times D_N}$ is defined element-wise by
 $[\otimes_{n=1}^N \mathbf{x}^{(n)}]_{d_1, \dots, d_N} = \prod_{n=1}^N x_{d_n}^{(n)}$, where $d_1 \in [D_1], \dots, d_N \in [D_N]$.

154 **Theorem 1.** Let $\mathcal{W}_{\text{TN}} \in \mathbb{R}^{D_1 \times \dots \times D_N}$ be a tensor generated by the locally connected tensor network
155 defined in Section 3.1, and let $\mathcal{A} \in \mathbb{R}^{D_1 \times \dots \times D_N}$. For any $\epsilon \in [0, \|\mathcal{A}\|/4]$, if $\|\mathcal{W}_{\text{TN}} - \mathcal{A}\| \leq \epsilon$,
156 then for all canonical partitions $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$ (Definition 2 it holds that $QE(\mathcal{A}; \mathcal{K}) \leq \ln(R) +$
157 $\frac{2\epsilon}{\|\mathcal{A}\|} \cdot \ln(D_{\mathcal{K}}) + 2\sqrt{\frac{2\epsilon}{\|\mathcal{A}\|}}$, where $D_{\mathcal{K}} := \min\{\prod_{n \in \mathcal{K}} D_n, \prod_{n \in \mathcal{K}^c} D_n\}$.³ In contrast, there exists
158 $\mathcal{A}' \in \mathbb{R}^{D_1 \times \dots \times D_N}$ such that for all canonical partitions $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$ it holds that $QE(\mathcal{A}'; \mathcal{K}) \geq$
159 $\min\{|\mathcal{K}|, |\mathcal{K}^c|\} \cdot \ln(\min_{n \in [N]} D_n)$.

160 *Proof sketch (proof in Appendix K.2).* In general, the entanglements that a tensor network supports
161 can be upper bounded through cuts in its graph [15, 35]. For the locally connected tensor network,
162 these bounds imply that $QE(\mathcal{W}_{\text{TN}}; \mathcal{K}) \leq \ln(R)$ for any canonical partition $(\mathcal{K}, \mathcal{K}^c)$. The first part of
163 the theorem then follows by showing that if \mathcal{W}_{TN} and \mathcal{A} are close, then so are their entanglements.
164 The second part is established using a construction from [17], providing a tensor with maximal
165 entanglements under all partitions of its axes. \square

166 **Theorem 2.** Let $\mathcal{A} \in \mathbb{R}^{D_1 \times \dots \times D_N}$ and $\epsilon > 0$. Suppose that for all canonical partitions $(\mathcal{K}, \mathcal{K}^c) \in$
167 \mathcal{C}_N (Definition 2) it holds that $QE(\mathcal{A}; \mathcal{K}) \leq \frac{\epsilon^2}{(2N-3)\|\mathcal{A}\|^2} \cdot \ln(R)$. Then, there exists an assignment
168 for the tensors constituting the locally connected tensor network (defined in Section 3.1) such that it
169 generates $\mathcal{W}_{\text{TN}} \in \mathbb{R}^{D_1 \times \dots \times D_N}$ satisfying $\|\mathcal{W}_{\text{TN}} - \mathcal{A}\| \leq \epsilon$.

170 *Proof sketch (proof in Appendix K.3).* We show that if \mathcal{A} has low entanglement under a canonical
171 partition $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$, then the singular values of $\llbracket \mathcal{A}; \mathcal{K} \rrbracket$ must decay rapidly (recall that $\llbracket \mathcal{A}; \mathcal{K} \rrbracket$ is
172 the arrangement of \mathcal{A} as a matrix where rows correspond to axes \mathcal{K} and columns correspond to the
173 remaining axes). The approximation guarantee is then obtained through a construction from [25],
174 which is based on truncated singular value decompositions of every $\llbracket \mathcal{A}; \mathcal{K} \rrbracket$ for $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$. \square

175 4 Low Entanglement Under Canonical Partitions Is Necessary and 176 Sufficient for Accurate Prediction

177 This section considers the locally connected tensor network from Section 3.1 in a machine learning
178 setting. We show that attaining low population loss amounts to fitting a tensor defined by the data
179 distribution, whose axes correspond to features (Section 4.1). Applying the theorems of Section 3.2,
180 we then conclude that the locally connected tensor network is capable of accurate prediction if and only
181 if the data distribution admits low entanglement under canonical partitions of features (Section 4.2).
182 This conclusion is corroborated through experiments, demonstrating that the performance of common
183 locally connected neural networks (including convolutional, recurrent, and local self-attention neural
184 networks) is inversely correlated with the entanglement under canonical partitions of features in
185 the data (Section 4.3). For conciseness, the treatment in this section is limited to one-dimensional
186 (sequential) models and data; see Appendix I.2 for an extension to arbitrary dimensions.

187 4.1 Accurate Prediction Is Equivalent to Fitting Data Tensor

188 As discussed in Section 3.1, the locally connected tensor network generating $\mathcal{W}_{\text{TN}} \in \mathbb{R}^{D_1 \times \dots \times D_N}$
189 is equivalent to a locally connected neural network, whose forward pass over a data instance
190 $(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)})$ yields $\langle \otimes_{n=1}^N \mathbf{x}^{(n)}, \mathcal{W}_{\text{TN}} \rangle$, where $\mathbf{x}^{(1)} \in \mathbb{R}^{D_1}, \dots, \mathbf{x}^{(N)} \in \mathbb{R}^{D_N}$. Motivated by this
191 fact, we consider a binary classification setting, in which the label y of the instance $(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)})$
192 is either 1 or -1 , and the prediction \hat{y} is taken to be the sign of the output of the neural network,
193 *i.e.* $\hat{y} = \text{sign}(\langle \otimes_{n=1}^N \mathbf{x}^{(n)}, \mathcal{W}_{\text{TN}} \rangle)$.

194 Suppose we are given a training set of labeled instances $\{((\mathbf{x}^{(1,m)}, \dots, \mathbf{x}^{(N,m)}), y^{(m)})\}_{m=1}^M$ drawn
195 i.i.d. from some distribution, and we would like to learn the parameters of the neural network through
196 the soft-margin support vector machine (SVM) objective, *i.e.* by optimizing:

$$\min_{\|\mathcal{W}_{\text{TN}}\| \leq B} \frac{1}{M} \sum_{m=1}^M \max\left\{0, 1 - y^{(m)} \langle \otimes_{n=1}^N \mathbf{x}^{(n,m)}, \mathcal{W}_{\text{TN}} \rangle\right\}, \quad (1)$$

197 for a predetermined constant $B > 0$. We assume instances are normalized, *i.e.* the distribution is
198 such that all vectors constituting an instance have norm no greater than one. We also assume that

³If $\mathcal{A} = 0$, then $\epsilon = 0$. In this case, the expression $\epsilon/\|\mathcal{A}\|$ is by convention equal to zero.

199 $B \leq 1$. In this case $|y^{(m)} \langle \otimes_{n=1}^N \mathbf{x}^{(n,m)}, \mathcal{W}_{\text{TN}} \rangle| \leq 1$, so our optimization problem can be expressed
 200 as $\min_{\|\mathcal{W}_{\text{TN}}\| \leq B} 1 - \langle \mathcal{D}_{\text{emp}}, \mathcal{W}_{\text{TN}} \rangle$, where

$$\mathcal{D}_{\text{emp}} := \frac{1}{M} \sum_{m=1}^M y^{(m)} \cdot \otimes_{n=1}^N \mathbf{x}^{(n,m)} \quad (2)$$

201 is referred to as the *empirical data tensor*. This means that the accuracy over the training data is
 202 determined by how large the inner product $\langle \mathcal{D}_{\text{emp}}, \mathcal{W}_{\text{TN}} \rangle$ is.

203 Disregarding the degenerate case of $\mathcal{D}_{\text{emp}} = 0$ (i.e. that in which the optimized ob-
 204 jective is constant), the inner products $\langle \mathcal{D}_{\text{emp}}, \mathcal{W}_{\text{TN}} \rangle$ and $\langle \mathcal{D}_{\text{emp}} / \|\mathcal{D}_{\text{emp}}\|, \mathcal{W}_{\text{TN}} \rangle$ differ by
 205 only a multiplicative (positive) constant, so fitting the training data amounts to optimizing
 206 $\max_{\|\mathcal{W}_{\text{TN}}\| \leq B} \langle \mathcal{D}_{\text{emp}} / \|\mathcal{D}_{\text{emp}}\|, \mathcal{W}_{\text{TN}} \rangle$. If \mathcal{W}_{TN} can represent some \mathcal{W} , then it can also represent $c \cdot$
 207 \mathcal{W} for every $c \in \mathbb{R}$. Thus, we may equivalently optimize $\max_{\mathcal{W}_{\text{TN}}} \langle \mathcal{D}_{\text{emp}} / \|\mathcal{D}_{\text{emp}}\|, \mathcal{W}_{\text{TN}} / \|\mathcal{W}_{\text{TN}}\| \rangle$
 208 and multiplying the result by B . Fitting the training data therefore boils down to minimizing
 209 $\left\| \frac{\mathcal{W}_{\text{TN}}}{\|\mathcal{W}_{\text{TN}}\|} - \frac{\mathcal{D}_{\text{emp}}}{\|\mathcal{D}_{\text{emp}}\|} \right\|$. In other words, the accuracy achievable over the training data is determined by
 210 the extent to which $\frac{\mathcal{W}_{\text{TN}}}{\|\mathcal{W}_{\text{TN}}\|}$ can fit the normalized empirical data tensor $\frac{\mathcal{D}_{\text{emp}}}{\|\mathcal{D}_{\text{emp}}\|}$.

211 The arguments above are independent of the training set size, and in fact apply to the population loss
 212 as well, in which case \mathcal{D}_{emp} is replaced by the *population data tensor*:

$$\mathcal{D}_{\text{pop}} := \mathbb{E}_{(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)})} [y \cdot \otimes_{n=1}^N \mathbf{x}^{(n)}]. \quad (3)$$

213 Disregarding the degenerate case of $\mathcal{D}_{\text{pop}} = 0$ (i.e. that in which the population loss is constant), it
 214 follows that the achievable accuracy over the population is determined by the extent to which $\frac{\mathcal{W}_{\text{TN}}}{\|\mathcal{W}_{\text{TN}}\|}$
 215 can fit the normalized population data tensor $\frac{\mathcal{D}_{\text{pop}}}{\|\mathcal{D}_{\text{pop}}\|}$. We refer to the minimal distance from it as the
 216 *suboptimality in achievable accuracy*.

217 **Definition 3.** In the context of the classification setting above, the *suboptimality in achievable*
 218 *accuracy* is $\text{SubOpt} := \min_{\mathcal{W}_{\text{TN}}} \left\| \frac{\mathcal{W}_{\text{TN}}}{\|\mathcal{W}_{\text{TN}}\|} - \frac{\mathcal{D}_{\text{pop}}}{\|\mathcal{D}_{\text{pop}}\|} \right\|$.

219 4.2 Necessary and Sufficient Condition for Accurate Prediction

220 In the classification setting of Section 4.1, by invoking Theorems 1 and 2 from Section 3.2, we
 221 conclude that the suboptimality in achievable accuracy is small if and only if the population data
 222 tensor \mathcal{D}_{pop} admits low entanglement under the canonical partitions of its axes (Definition 2).

223 **Corollary 1.** Consider the classification setting of Section 4.1, and let $\epsilon \in [0, 1/4]$. If there exists
 224 a canonical partition $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$ (Definition 2) under which $QE(\mathcal{D}_{\text{pop}}; \mathcal{K}) > \ln(R) + 2\epsilon \cdot$
 225 $\ln(D_{\mathcal{K}}) + 2\sqrt{2}\epsilon$, where $D_{\mathcal{K}} := \min\{\prod_{n \in \mathcal{K}} D_n, \prod_{n \in \mathcal{K}^c} D_n\}$, then $\text{SubOpt} > \epsilon$. Conversely, if for
 226 all $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$ it holds that $QE(\mathcal{D}_{\text{pop}}; \mathcal{K}) \leq \frac{\epsilon^2}{8N-12} \cdot \ln(R)$, then $\text{SubOpt} \leq \epsilon$.

227 *Proof sketch (proof in Appendix K.5).* Follows from Theorems 1 and 2 after accounting for the nor-
 228 malization of \mathcal{W}_{TN} in the suboptimality in achievable accuracy. \square

229 Directly evaluating the conditions required by Corollary 1 — low entanglement under canonical
 230 partitions for \mathcal{D}_{pop} — is impractical, since: (i) \mathcal{D}_{pop} is defined via an unknown data distribution
 231 (Equation (3)); and (ii) computing the entanglements involves taking singular value decompositions
 232 of matrices with size exponential in the number of input variables N . Fortunately, as Proposition 1
 233 (in Appendix C) shows, the entanglements of \mathcal{D}_{pop} under all partitions are with high probability well-
 234 approximated by the entanglements of the empirical data tensor \mathcal{D}_{emp} . Moreover, the entanglement
 235 of \mathcal{D}_{emp} under any partition can be computed efficiently, without explicitly storing or manipulating
 236 an exponentially large matrix — see Appendix E for an algorithm (originally proposed in [40]).
 237 Overall, we obtain an efficiently computable criterion (low entanglement under canonical partitions
 238 for \mathcal{D}_{emp}), that with high probability is both necessary and sufficient for low suboptimality in
 239 achievable accuracy — see Corollary 2 (in Appendix C).

240 4.3 Empirical Demonstration

241 Corollary 2 establishes that, with high probability, the locally connected tensor network (from Sec-
 242 tion 3.1) can achieve high prediction accuracy if and only if the empirical data tensor (Equation (2))
 243 admits low entanglement under canonical partitions of its axes. We corroborate our formal analysis

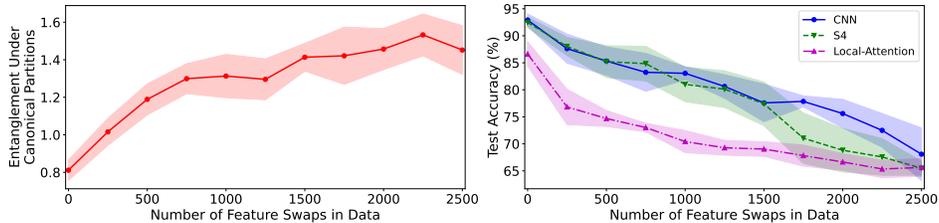


Figure 3: The prediction accuracies of common locally connected neural networks are inversely correlated with the entanglements of the data under canonical partitions of features, in compliance with our theory (Sections 4.1 and 4.2). **Left:** Average entanglement under canonical partitions (Definition 2) of the empirical data tensor (Equation (2)), for binary classification variants of the Speech Commands audio dataset [62] obtained by performing random position swaps between features. **Right:** Test accuracies achieved by a convolutional neural network (CNN) [16], S4 (a popular class of recurrent neural networks; see [26]), and a local self-attention model [46], against the number of random feature swaps performed to create the dataset. **All:** Reported are the means and standard deviations of the quantities specified above, taken over ten different random seeds. See Appendix I.2.3 for experiments over (two-dimensional) image data and Appendix J.2 for implementation details.

244 through experiments, demonstrating that its conclusions carry over to common locally connected
 245 architectures. Namely, applying convolutional neural networks, S4 (a popular recurrent neural
 246 network; see [26]), and a local self-attention model [46] to different datasets, we show that the
 247 achieved test accuracy is inversely correlated with the entanglements of the empirical data tensor
 248 under canonical partitions. Below is a description of experiments with one-dimensional (*i.e.* sequential)
 249 models and data. Additional experiments with two-dimensional (imagery) models and data are given
 250 in Appendix I.2.3.

251 Discerning the relation between entanglements of the empirical data tensor and performance (pre-
 252 diction accuracy) of locally connected neural networks requires datasets admitting different entan-
 253 glements. A potential way to acquire such datasets is as follows. First, select one a dataset which
 254 locally connected neural networks perform well, in the hopes that it admits low entanglement under
 255 canonical partitions; natural candidates are datasets comprising images, text or audio. Subsequently,
 256 create “shuffled” variants of the dataset by repeatedly swapping the position of two features chosen at
 257 random. This erodes the original arrangement of features in the data, and is expected to yield higher
 258 entanglement under canonical partitions.

259 We followed the blueprint above for a binary classification version of the Speech Commands audio
 260 dataset [62]. Figure 3 presents test accuracies achieved by a convolutional neural network, S4,
 261 and a local self-attention model, as well as average entanglement under canonical partitions of the
 262 empirical data tensor, against the number of random feature swaps performed to create the dataset. As
 263 expected, when the number of swaps increases, the average entanglement under canonical partitions
 264 becomes higher. At the same time, in accordance with our theory, the prediction accuracies of the
 265 locally connected neural networks substantially deteriorate, showing an inverse correlation with the
 266 entanglement under canonical partitions.

267 5 Enhancing Suitability of Data to Locally Connected Neural Networks

268 Our analysis (Sections 3 and 4) suggests that a data distribution is suitable for locally connected neural
 269 networks if and only if it admits low entanglement under canonical partitions of features. Motivated
 270 by this observation, we derive a preprocessing algorithm aimed to enhance the suitability of a data
 271 distribution to locally connected neural networks (Section 5.1 and Appendix F). Empirical evaluations
 272 demonstrate that it significantly improves prediction accuracies of common locally connected neural
 273 networks on various datasets (Section 5.2). For conciseness, the treatment in this section is limited
 274 to one-dimensional (sequential) models and data; see Appendix I.3 for an extension to arbitrary
 275 dimensions.

276 5.1 Search for Feature Arrangement With Low Entanglement Under Canonical Partitions

277 Our analysis naturally leads to a recipe for enhancing the suitability of a data distribution to locally
 278 connected neural networks: given a dataset, search for an arrangement of features which leads to low
 279 entanglement under canonical partitions, and then arrange the features accordingly. Formally, suppose
 280 we have $M \in \mathbb{N}$ training instances $\{((\mathbf{x}^{(1,m)}, \dots, \mathbf{x}^{(N,m)}), y^{(m)})\}_{m=1}^M$, where $y^{(m)} \in \{1, -1\}$ and
 281 $\mathbf{x}^{(n,m)} \in \mathbb{R}^D$ for $n \in [N], m \in [M]$, with $D \in \mathbb{N}$. Assume without loss of generality that N is a
 282 power of two (if this is not the case we may add constant features as needed). The aforementioned

283 recipe boils down to a search for a permutation $\pi : [N] \rightarrow [N]$, which when applied to feature indices
 284 leads the empirical data tensor \mathcal{D}_{emp} (Equation (2)) to admit low entanglement under the canonical
 285 partitions of its axes (Definition 2).

286 A greedy realization of the foregoing search is as follows. Initially, partition the features into two
 287 equally sized sets $\mathcal{K}_{1,1} \subset [N]$ and $\mathcal{K}_{1,2} := [N] \setminus \mathcal{K}_{1,1}$ such that the entanglement of \mathcal{D}_{emp} with
 288 respect to $(\mathcal{K}_{1,1}, \mathcal{K}_{1,2})$ is minimal. That is, find $\mathcal{K}_{1,1} \in \operatorname{argmin}_{\mathcal{K} \subset [N], |\mathcal{K}|=N/2} QE(\mathcal{D}_{\text{emp}}; \mathcal{K})$. The
 289 permutation π will map $\mathcal{K}_{1,1}$ to coordinates $\{1, \dots, \frac{N}{2}\}$ and $\mathcal{K}_{1,2}$ to $\{\frac{N}{2} + 1, \dots, N\}$. Then, partition
 290 $\mathcal{K}_{1,1}$ into two equally sized sets $\mathcal{K}_{2,1} \subset \mathcal{K}_{1,1}$ and $\mathcal{K}_{2,2} := \mathcal{K}_{1,1} \setminus \mathcal{K}_{2,1}$ such that the average of entan-
 291 glements induced by these sets is minimal, *i.e.* $\mathcal{K}_{2,1} \in \operatorname{argmin}_{\mathcal{K} \subset \mathcal{K}_{1,1}, |\mathcal{K}|=|\mathcal{K}_{1,1}|/2} \frac{1}{2} [QE(\mathcal{D}_{\text{emp}}; \mathcal{K}) +$
 292 $QE(\mathcal{D}_{\text{emp}}; \mathcal{K}_{1,1} \setminus \mathcal{K})]$. The permutation π will map $\mathcal{K}_{2,1}$ to coordinates $\{1, \dots, \frac{N}{4}\}$ and $\mathcal{K}_{2,2}$ to
 293 $\{\frac{N}{4} + 1, \dots, \frac{N}{2}\}$. A partition of $\mathcal{K}_{1,2}$ into two equally sized sets $\mathcal{K}_{2,3}$ and $\mathcal{K}_{2,4}$ is obtained similarly,
 294 where π will map $\mathcal{K}_{2,3}$ to coordinates $\{\frac{N}{2} + 1, \dots, \frac{3N}{4}\}$ and $\mathcal{K}_{2,4}$ to $\{\frac{3N}{4} + 1, \dots, N\}$. Continuing
 295 in the same fashion, until we reach subsets $\mathcal{K}_{L,1}, \dots, \mathcal{K}_{L,N}$ consisting of a single feature index each,
 296 fully specifies the permutation π .

297 Unfortunately, the step lying at the heart of the above scheme — finding a balanced partition that
 298 minimizes average entanglement — is computationally prohibitive, and we are not aware of any tools
 299 that alleviate the computational difficulty. However, as discussed in Appendix F, if one replaces
 300 entanglement with an appropriate surrogate measure, then each search for a balanced partition
 301 minimizing average entanglement converts into a *minimum balanced cut problem*, which enjoys a
 302 wide array of established approximation tools [29]. We thus obtain a practical algorithm for enhancing
 303 the suitability of a data distribution to locally connected neural networks.

304 5.2 Experiments

305 We empirically evaluate our feature rearrangement method (detailed in Appendix F) using common
 306 locally connected neural networks — a convolutional neural network, an S4 (popular recurrent neural
 307 network; see [26]), and a local self-attention model [46] — over randomly permuted audio datasets
 308 (Section 5.2.1) and several tabular datasets (Section 5.2.2). For brevity, we defer some implementation
 309 details and experiments to Appendix J. Additional experiments with two-dimensional data are given
 310 in Appendix I.3.3.

311 5.2.1 Randomly Permuted Audio Datasets

312 Section 4.3 demonstrated that audio data admits low entanglement under canonical partitions of
 313 features, and that randomly permuting the position of features leads this entanglement to increase,
 314 while substantially degrading the prediction accuracy of locally connected neural networks. A
 315 sensible test for our method is to evaluate its ability to recover performance lost due to the random
 316 permutation of features.

317 For the Speech Commands dataset [62], Table 1 compares prediction accuracies of locally connected
 318 neural networks on the data: (i) subject to a random permutation of features; (ii) attained after
 319 rearranging the randomly permuted features via our method; and (iii) attained after rearranging the
 320 randomly permuted features via IGTD [65] — a heuristic scheme designed for convolutional neural
 321 networks (see Appendix A). As can be seen, our method leads to significant improvements, surpassing
 322 those brought forth by IGTD. Note that the performance lost due to the random permutation of features
 323 is not entirely recovered.⁴ We believe this relates to phenomena outside the scope of the theory
 324 underlying our method (Sections 3 and 4), for example translation invariance in data being beneficial
 325 in terms of generalization. Investigation of such phenomena and suitable modification of our method
 326 are regarded as promising directions for future work.

327 5.2.2 Tabular Datasets

328 The prediction accuracies of locally connected neural networks on tabular data, *i.e.* on data in
 329 which features are arranged arbitrarily, is known to be subpar [55]. Tables 2 and 5 report results
 330 of experiments with locally connected neural networks over standard tabular benchmarks (namely
 331 “semeion”, “isolet” and “dna” [59]), demonstrating that arranging features via our method leads to
 332 significant improvements in prediction accuracies, surpassing improvements brought forth by IGTD
 333 (a heuristic scheme designed for convolutional neural networks [65]). Note that our method does
 334 not lead to state of the art prediction accuracies on the evaluated benchmarks.⁵ However, the results

⁴Accuracies on the original data are 59.8, 69.6 and 48.1 for CNN, S4 and Local-Attention, respectively.

⁵XGBoost, *e.g.*, achieves prediction accuracies 91, 95.2 and 96 over semeion, isolet and dna, respectively.

Table 1: Arranging features of randomly permuted audio data via our method (detailed in Appendix F) significantly improves the prediction accuracies of locally connected neural networks. Reported are test accuracies (mean and standard deviation over ten random seeds) of a convolutional neural network (CNN), S4 (a popular recurrent neural network; see [26]), and a local self-attention model [46], over the Speech Commands dataset [62] subject to different arrangements of features: (i) a random arrangement; (ii) an arrangement provided by applying our method to the random arrangement; and (iii) an arrangement provided by applying an adaptation of IGTD [65] — a heuristic scheme designed for convolutional neural networks — to the random arrangement. For each model, we highlight (in boldface) the highest mean accuracy if the difference between that and the second-highest mean accuracy is statistically significant (namely, is larger than the standard deviation corresponding to the former). Our method leads to significant improvements in prediction accuracies, surpassing the improvements brought forth by IGTD. See Appendix J for experiments demonstrating its scalability and implementation details.

	Randomly Permuted	Our Method	IGTD
CNN	5.2 ± 0.7	17.4 ± 1.7	6.1 ± 0.4
S4	9.5 ± 0.6	30.3 ± 1.6	13 ± 2.4
Local-Attention	7.8 ± 0.3	12.9 ± 0.7	6.4 ± 0.5

Table 2: Arranging features of tabular datasets via our method (detailed in Appendix F) significantly improves the prediction accuracies of locally connected neural networks. Reported are results of experiments analogous to those of Table 1, but with the “semeion” and “isolet” tabular classification datasets [59]. Since to the arrangement of features in a tabular dataset is intended to be arbitrary, we regard as a baseline the prediction accuracies attained with a random permutation of features. As in the experiment of Table 1, rearranging the features according to our method leads to significant improvements in prediction accuracies, surpassing the improvements brought forth by IGTD. See Appendix J for experiments with an additional tabular dataset (“dna”) and implementation details.

	Dataset: semeion			Dataset: isolet		
	Baseline	Our Method	IGTD	Baseline	Our Method	IGTD
CNN	77.7 ± 1.4	80.0 ± 1.8	78.9 ± 1.9	91.0 ± 0.6	92.5 ± 0.4	92.0 ± 0.6
S4	82.5 ± 1.1	89.7 ± 0.5	86.0 ± 0.7	92.3 ± 0.4	93.4 ± 0.3	92.7 ± 0.5
Local-Attention	60.9 ± 4.9	78.0 ± 1.7	67.8 ± 2.6	82.0 ± 1.6	89.0 ± 0.6	85.7 ± 1.9

335 suggest that it renders locally connected neural networks a viable option for tabular data. This option
 336 is particularly appealing in when the number of features is large settings, where many alternative
 337 approaches (*e.g.* ones involving fully connected neural networks) are impractical.

338 6 Conclusion

339 The question of what makes a data distribution suitable for deep learning is a fundamental open
 340 problem. Focusing on locally connected neural networks — a prevalent family of deep learning
 341 architectures that includes as special cases convolutional neural networks, recurrent neural networks
 342 (in particular the recent S4 models) and local self-attention models — we address this problem by
 343 adopting theoretical tools from quantum physics. Our main theoretical result states that a certain
 344 locally connected neural network is capable of accurate prediction (*i.e.* can express a solution with
 345 low population loss) over a data distribution *if and only if* the data distribution admits low quantum
 346 entanglement under certain canonical partitions of features. Experiments with widespread locally
 347 connected neural networks corroborate this finding.

348 Our theory suggests that the suitability of a data distribution to locally connected neural networks
 349 may be enhanced by arranging features such that low entanglement under canonical partitions is
 350 attained. Employing a certain surrogate for entanglement, we show that this arrangement can be
 351 implemented efficiently, and that it leads to substantial improvements in the prediction accuracies of
 352 common locally connected neural networks on various datasets.

353 The data modalities to which deep learning is most commonly applied — namely ones involving
 354 images, text and audio — are often regarded as natural (as opposed to, for example, tabular data
 355 fusing heterogeneous information). We believe the difficulty in explaining the suitability of such
 356 modalities to deep learning may be due to a shortage in tools for formally reasoning about natural
 357 data. Concepts and tools from physics — a branch of science concerned with formally reasoning
 358 about natural phenomena — may be key to overcoming said difficulty. We hope that our use of
 359 quantum entanglement will encourage further research along this line.

360 References

- 361 [1] Emmanuel Abbe and Colin Sandon. Provable limitations of deep learning. *arXiv preprint arXiv:1812.06369*,
362 2018.
- 363 [2] Emilio Rafael Balda, Arash Behboodi, and Rudolf Mathar. A tensor analysis on dense connectivity via
364 convolutional arithmetic circuits. *Preprint*, 2018.
- 365 [3] Rajendra Bhatia. *Matrix analysis*. Springer-Verlag, 1997.
- 366 [4] Alon Brutzkus and Amir Globerson. Globally optimal gradient descent for a convnet with gaussian inputs.
367 In *International Conference on Machine Learning (ICML)*, pages 605–614, 2017.
- 368 [5] Alon Brutzkus, Amir Globerson, Eran Malach, Alon Regev Netser, and Shai Shalev-Schwartz. Efficient
369 learning of CNNs using patch based features. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba
370 Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on*
371 *Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 2336–2356. PMLR,
372 17–23 Jul 2022. URL <https://proceedings.mlr.press/v162/brutzkus22a.html>.
- 373 [6] Song Cheng, Jing Chen, and Lei Wang. Information perspective to probabilistic modeling: Boltzmann
374 machines versus born machines. *Entropy*, 20(8):583, 2018.
- 375 [7] Nadav Cohen and Amnon Shashua. Simnets: A generalization of convolutional networks. *Advances in*
376 *Neural Information Processing Systems (NeurIPS), Deep Learning Workshop*, 2014.
- 377 [8] Nadav Cohen and Amnon Shashua. Convolutional rectifier networks as generalized tensor decompositions.
378 *International Conference on Machine Learning (ICML)*, 2016.
- 379 [9] Nadav Cohen and Amnon Shashua. Inductive bias of deep convolutional networks through pooling
380 geometry. *International Conference on Learning Representations (ICLR)*, 2017.
- 381 [10] Nadav Cohen, Or Sharir, and Amnon Shashua. Deep simnets. *IEEE Conference on Computer Vision and*
382 *Pattern Recognition (CVPR)*, 2016.
- 383 [11] Nadav Cohen, Or Sharir, and Amnon Shashua. On the expressive power of deep learning: A tensor analysis.
384 *Conference On Learning Theory (COLT)*, 2016.
- 385 [12] Nadav Cohen, Or Sharir, Yoav Levine, Ronen Tamari, David Yakira, and Amnon Shashua. Analysis and
386 design of convolutional networks via hierarchical tensor decompositions. *Intel Collaborative Research*
387 *Institute for Computational Intelligence (ICRI-CI) Special Issue on Deep Learning Theory*, 2017.
- 388 [13] Nadav Cohen, Ronen Tamari, and Amnon Shashua. Boosting dilated convolutional networks with mixed
389 tensor decompositions. *International Conference on Learning Representations (ICLR)*, 2018.
- 390 [14] Ian Convy, William Huggins, Haoran Liao, and K Birgitta Whaley. Mutual information scaling for tensor
391 network machine learning. *Machine learning: science and technology*, 3(1):015017, 2022.
- 392 [15] Shawn X Cui, Michael H Freedman, Or Sattath, Richard Stong, and Greg Minton. Quantum max-flow/min-
393 cut. *Journal of Mathematical Physics*, 57(6):062206, 2016.
- 394 [16] Wei Dai, Chia Dai, Shuhui Qu, Juncheng Li, and Samarjit Das. Very deep convolutional neural networks
395 for raw waveforms. In *2017 IEEE international conference on acoustics, speech and signal processing*
396 *(ICASSP)*, pages 421–425. IEEE, 2017.
- 397 [17] Dong-Ling Deng, Xiaopeng Li, and S Das Sarma. Quantum entanglement in neural network states.
398 *Physical Review X*, 7(2):021021, 2017.
- 399 [18] Persi Diaconis and Mehrdad Shahshahani. Generating a random permutation with random transpositions.
400 *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 57(2):159–179, 1981.
- 401 [19] Simon S Du and Surbhi Goel. Improved learning of one-hidden-layer convolutional neural networks with
402 overlaps. *arXiv preprint arXiv:1805.07798*, 2018.
- 403 [20] Simon S Du, Jason D Lee, Yuandong Tian, Aarti Singh, and Barnabas Poczos. Gradient descent learns
404 one-hidden-layer cnn: Don’t be afraid of spurious local minima. In *International Conference on Machine*
405 *Learning (ICML)*, pages 1339–1348, 2018.
- 406 [21] Anatoly Dymarsky and Kirill Pavlenko. Tensor network to learn the wave function of data. *Physical*
407 *Review Research*, 4(4):043111, 2022.

- 408 [22] Timo Felser, Marco Trenti, Lorenzo Sestini, Alessio Gianelle, Davide Zuliani, Donatella Lucchesi, and
409 Simone Montangero. Quantum-inspired machine learning on high-energy physics data. *npj Quantum*
410 *Information*, 7(1):1–8, 2021.
- 411 [23] Michael R Garey and David S Johnson. Computers and intractability. *A Guide to the Theory of NP-*
412 *Completeness*, 1979.
- 413 [24] Rong Ge, Yunwei Ren, Xiang Wang, and Mo Zhou. Understanding deflation process in over-parametrized
414 tensor decomposition. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- 415 [25] Lars Grasedyck. Hierarchical singular value decomposition of tensors. *SIAM journal on matrix analysis*
416 *and applications*, 31(4):2029–2054, 2010.
- 417 [26] Albert Gu, Karan Goel, and Christopher Ré. Efficiently modeling long sequences with structured state
418 spaces. *International Conference on Learning Representations (ICLR)*, 2022.
- 419 [27] Siu-Wai Ho and Raymond W Yeung. The interplay between entropy and variational distance. *IEEE*
420 *Transactions on Information Theory*, 56(12):5906–5929, 2010.
- 421 [28] Zhih-Ahn Jia, Lu Wei, Yu-Chun Wu, Guang-Can Guo, and Guo-Ping Guo. Entanglement area law for
422 shallow and deep quantum neural network states. *New Journal of Physics*, 22(5):053022, 2020.
- 423 [29] George Karypis and Vipin Kumar. A fast and high quality multilevel scheme for partitioning irregular
424 graphs. *SIAM Journal on scientific Computing*, 20(1):359–392, 1998.
- 425 [30] Valentin Khrulkov, Alexander Novikov, and Ivan Oseledets. Expressive power of recurrent neural networks.
426 *International Conference on Learning Representations (ICLR)*, 2018.
- 427 [31] Valentin Khrulkov, Oleksii Hrinchuk, and Ivan Oseledets. Generalized tensor models for recurrent neural
428 networks. *International Conference on Learning Representations (ICLR)*, 2019.
- 429 [32] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint*
430 *arXiv:1412.6980*, 2014.
- 431 [33] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.
- 432 [34] Yoav Levine, Or Sharir, and Amnon Shashua. Benefits of depth for long-term memory of recurrent
433 networks. *International Conference on Learning Representations (ICLR) Workshop*, 2018.
- 434 [35] Yoav Levine, David Yakira, Nadav Cohen, and Amnon Shashua. Deep learning and quantum entanglement:
435 Fundamental connections with implications to network design. *International Conference on Learning*
436 *Representations (ICLR)*, 2018.
- 437 [36] Yoav Levine, Or Sharir, Nadav Cohen, and Amnon Shashua. Quantum entanglement in deep learning
438 architectures. *To appear in Physical Review Letters*, 2019.
- 439 [37] Jingling Li, Yanchao Sun, Jiahao Su, Taiji Suzuki, and Furong Huang. Understanding generalization in
440 deep learning via tensor methods. In *International Conference on Artificial Intelligence and Statistics*,
441 pages 504–515. PMLR, 2020.
- 442 [38] Sirui Lu, Márton Kanász-Nagy, Ivan Kukuljan, and J Ignacio Cirac. Tensor networks and efficient
443 descriptions of classical data. *arXiv preprint arXiv:2103.06872*, 2021.
- 444 [39] Eran Malach and Shai Shalev-Shwartz. A provably correct algorithm for deep learning that actually works.
445 *arXiv preprint arXiv:1803.09522*, 2018.
- 446 [40] John Martyn, Guifre Vidal, Chase Roberts, and Stefan Leichenauer. Entanglement and tensor networks for
447 supervised image classification. *arXiv preprint arXiv:2007.06082*, 2020.
- 448 [41] Kento Nozawa and Issei Sato. Empirical evaluation and theoretical analysis for representation learning: A
449 survey. *arXiv preprint arXiv:2204.08226*, 2022.
- 450 [42] Román Orús. A practical introduction to tensor networks: Matrix product states and projected entangled
451 pair states. *Annals of physics*, 349:117–158, 2014.
- 452 [43] Samet Oymak and Mahdi Soltanolkotabi. End-to-end learning of a convolutional neural network via deep
453 tensor decomposition. *arXiv preprint arXiv:1805.06523*, 2018.
- 454 [44] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming
455 Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NIPS-W*,
456 2017.

- 457 [45] Giovanni Puccetti. Measuring linear correlation between random vectors. *Available at SSRN 3116066*,
458 2022.
- 459 [46] Jack Rae and Ali Razavi. Do transformers need deep long-range memory? In *Proceedings of the 58th*
460 *Annual Meeting of the Association for Computational Linguistics*, Online, July 2020. Association for
461 Computational Linguistics. URL <https://www.aclweb.org/anthology/2020.acl-main.672>.
- 462 [47] Noam Razin and Nadav Cohen. Implicit regularization in deep learning may not be explainable by norms.
463 In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- 464 [48] Noam Razin, Asaf Maman, and Nadav Cohen. Implicit regularization in tensor factorization. *International*
465 *Conference on Machine Learning (ICML)*, 2021.
- 466 [49] Noam Razin, Asaf Maman, and Nadav Cohen. Implicit regularization in hierarchical tensor factorization
467 and deep convolutional neural networks. *International Conference on Machine Learning (ICML)*, 2022.
- 468 [50] Noam Razin, Tom Verbin, and Nadav Cohen. On the ability of graph neural networks to model interactions
469 between vertices. *arXiv preprint arXiv:2211.16494*, 2022.
- 470 [51] Lorenzo Rosasco, Mikhail Belkin, and Ernesto De Vito. On learning with integral operators. *J. Mach.*
471 *Learn. Res.*, 11:905–934, mar 2010. ISSN 1532-4435.
- 472 [52] Shai Shalev-Shwartz, Ohad Shamir, and Shaked Shammah. Failures of gradient-based deep learning. In
473 *International Conference on Machine Learning*, pages 3067–3075. PMLR, 2017.
- 474 [53] Or Sharir and Amnon Shashua. On the expressive power of overlapping architectures of deep learning.
475 *International Conference on Learning Representations (ICLR)*, 2018.
- 476 [54] Or Sharir, Ronen Tamari, Nadav Cohen, and Amnon Shashua. Tensorial mixture models. *arXiv preprint*
477 *arXiv:1610.04167*, 2016.
- 478 [55] Ravid Shwartz-Ziv and Amitai Armon. Tabular data: Deep learning is not all you need. *Information*
479 *Fusion*, 81:84–90, 2022.
- 480 [56] Daniel A Spielman and Shang-Hua Teng. Spectral sparsification of graphs. *SIAM Journal on Computing*,
481 40(4):981–1025, 2011.
- 482 [57] E Miles Stoudenmire. Learning relevant features of data with multi-scale tensor networks. *Quantum*
483 *Science and Technology*, 3(3):034003, 2018.
- 484 [58] Edwin Stoudenmire and David J Schwab. Supervised learning with tensor networks. In *Advances in Neural*
485 *Information Processing Systems (NeurIPS)*, 2016.
- 486 [59] Joaquin Vanschoren, Jan N. van Rijn, Bernd Bischl, and Luis Torgo. Openml: networked science in
487 machine learning. *SIGKDD Explorations*, 15(2):49–60, 2013. doi: 10.1145/2641190.2641198. URL
488 <http://doi.acm.org/10.1145/2641190.264119>.
- 489 [60] Qingcan Wang et al. Exponential convergence of the deep neural network approximation for analytic
490 functions. *arXiv preprint arXiv:1807.00297*, 2018.
- 491 [61] Xiang Wang, Chenwei Wu, Jason D Lee, Tengyu Ma, and Rong Ge. Beyond lazy training for over-
492 parameterized tensor decomposition. *Advances in Neural Information Processing Systems*, 33:21934–
493 21944, 2020.
- 494 [62] Pete Warden. Speech commands: A dataset for limited-vocabulary speech recognition. *arXiv preprint*
495 *arXiv:1804.03209*, 2018.
- 496 [63] Dmitry Yarotsky. Error bounds for approximations with deep relu networks. *Neural Networks*, 94:103–114,
497 2017.
- 498 [64] Ya-Hui Zhang. Entanglement entropy of target functions for image classification and convolutional neural
499 network. *arXiv preprint arXiv:1710.05520*, 2017.
- 500 [65] Yitan Zhu, Thomas Brettin, Fangfang Xia, Alexander Partin, Maulik Shukla, Hyunseung Yoo, Yvonne A
501 Evrard, James H Doroshov, and Rick L Stevens. Converting tabular data into images for deep learning
502 with convolutional neural networks. *Scientific reports*, 11(1):1–11, 2021.

503 A Related Work

504 Characterizing formal properties of data distributions that make them suitable for deep learning
 505 is a major open problem in the field. A number of papers provide sufficient conditions on a data
 506 distribution which imply that it is learnable by certain neural networks [4, 63, 39, 19, 20, 43, 60, 5].
 507 However, these sufficient conditions are restrictive, and are not argued to be necessary for any aspect
 508 of learning (*i.e.* for expressiveness, optimization or generalization). To the best of our knowledge,
 509 this paper is the first to derive a verifiable condition on a data distribution that is both necessary and
 510 sufficient for aspects of learning to be achievable by a neural network. We note that the condition we
 511 derive resembles past attempts to quantify the structure of data via quantum entanglement and mutual
 512 information [40, 14, 38, 6, 64, 28, 21]. However, such quantifications have not been formally related
 513 to learnability by neural networks.

514 The current paper follows a long line of research employing tensor networks as theoretical models
 515 for studying deep learning. This line includes works analyzing the expressiveness of different neural
 516 network architectures [11, 54, 8, 9, 53, 13, 34, 2, 35, 30, 36, 31, 50], their generalization [37], and
 517 the implicit regularization induced by their optimization [47, 48, 49, 61, 24]. Similarly to prior works
 518 we focus on expressiveness, yet our approach differs in that we incorporate the data distribution into
 519 the analysis and tackle the question of what makes data suitable for deep learning.

520 The algorithm we propose for enhancing the suitability of data to locally connected neural networks
 521 can be considered a form of representation learning. Representation learning is a vast field, far too
 522 broad to survey here (for an overview see [41]). Our algorithm, which learns a representation via
 523 rearrangement of features in the data, is complementary to most representation learning methods
 524 in the literature. A notable method that is also based on feature rearrangement is IGTD [65] — a
 525 heuristic scheme designed for convolutional neural networks. In contrast to IGTD, our algorithm is
 526 theoretically grounded. Moreover, we demonstrate empirically in Section 5 that it leads to higher
 527 prediction accuracies.

528 B Illustration of Canonical Partitions

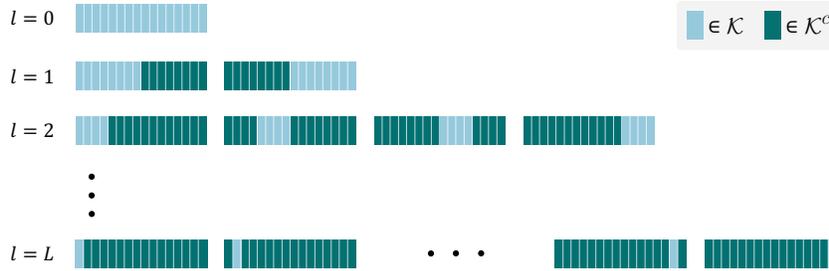


Figure 4: The canonical partitions of $[N]$, for $N = 2^L$ with $L \in \mathbb{N}$. Every $l \in \{0, \dots, L\}$ contributes 2^l canonical partitions, the n 'th one induced by $\mathcal{K} = \{2^{L-l} \cdot (n-1) + 1, \dots, 2^{L-l} \cdot n\}$.

529 C Efficiently Computable Criterion for Low Suboptimality in Achievable 530 Accuracy

531 In this appendix, we provide the formal claims deferred from Section 4.2. Specifically, Proposition 1
 532 shows that the entanglements of \mathcal{D}_{pop} under all partitions are with high probability well-approximated
 533 by the entanglements of the empirical data tensor \mathcal{D}_{emp} . Corollary 2 establishes the efficiently com-
 534 putable criterion (low entanglement under canonical partitions for \mathcal{D}_{emp}), that with high probability
 535 is both necessary and sufficient for low suboptimality in achievable accuracy.

536 **Proposition 1.** Consider the classification setting of Section 4.1, and let $\delta \in (0, 1)$ and
 537 $\gamma > 0$. If the training set size M satisfies $M \geq \frac{128 \ln(\frac{2}{\delta}) (\ln(\max_{\mathcal{K}' \subseteq [N]} \mathcal{D}_{\mathcal{K}'})^4)}{\|\mathcal{D}_{\text{pop}}\|^2 \gamma^4}$, where $\mathcal{D}_{\mathcal{K}'} :=$
 538 $\min\{\prod_{n \in \mathcal{K}'} D_n, \prod_{n \in \mathcal{K}'^c} D_n\}$ for $\mathcal{K}' \subseteq [N]$, then with probability at least $1 - \delta$:

$$|QE(\mathcal{D}_{\text{emp}}; \mathcal{K}) - QE(\mathcal{D}_{\text{pop}}; \mathcal{K})| \leq \gamma \text{ for all } \mathcal{K} \subseteq [N].$$

539 *Proof sketch (proof in Appendix K.4).* A standard generalization of the Hoeffding inequality to ran-
 540 dom vectors in a Hilbert space allows bounding $\|\mathcal{D}_{\text{emp}} - \mathcal{D}_{\text{pop}}\|$ with high probability. Bounding
 541 differences between entanglements via Euclidean distance concludes the proof. \square

542 **Corollary 2.** Consider the setting and notation of Corollary 1, with $\epsilon \neq 0$. For $\delta \in (0, 1)$, suppose
 543 that the training set size M satisfies $M \geq \frac{128 \ln(\frac{2}{\delta})((16N-24) \ln(\max_{\mathcal{K}' \in \mathcal{C}_N} D_{\mathcal{K}'})^4)}{\|\mathcal{D}_{\text{pop}}\|^2 (\ln(R) \cdot \epsilon^2)^4}$. Then, with prob-
 544 ability at least $1 - \delta$ the following hold. First, if there exists a canonical partition $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$
 545 (Definition 2) under which $QE(\mathcal{D}_{\text{emp}}; \mathcal{K}) > (1 + \frac{\epsilon^2}{16N-24}) \cdot \ln(R) + 2\epsilon \cdot \ln(D_{\mathcal{K}}) + 2\sqrt{2}\epsilon$, then:

$$\text{SubOpt} > \epsilon.$$

546 Second, if for all $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$ it holds that $QE(\mathcal{D}_{\text{emp}}; \mathcal{K}) \leq \frac{\epsilon^2}{16N-24} \cdot \ln(R)$, then:

$$\text{SubOpt} \leq \epsilon.$$

547 Moreover, the conditions above on the entanglements of \mathcal{D}_{emp} can be evaluated efficiently (in
 548 $\mathcal{O}(DN^2M^2 + NM^3)$ time $\mathcal{O}(DNM + M^2)$ and memory, where $D := \max_{n \in [N]} D_n$).

549 *Proof.* Implied by Corollary 1, Proposition 1 with $\gamma = \frac{\epsilon^2}{16N-24} \cdot \ln(R)$ and Algorithm 1 in Appendix E.
 550 \square

551 D Impossibility Result for Improving the Sufficient Condition in Theorem 2

552 The sufficient condition in Theorem 2 (from Section 3.2) for approximating $\mathcal{A} \in \mathbb{R}^{D_1 \times \dots \times D_N}$
 553 requires entanglements to approach zero as the desired approximation error ϵ does, in contrast to
 554 the necessary condition in Theorem 1 where they approach $\ln(R)$. As Proposition 2 shows, this is
 555 unavoidable in the absence of further knowledge regarding \mathcal{A} . However, if for all canonical partitions
 556 $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$ the singular values of $\llbracket \mathcal{A}; \mathcal{K} \rrbracket$ trailing after the R 'th one are small, then we can also
 557 guarantee an assignment for the locally connected tensor network satisfying $\|\mathcal{W}_{\text{TN}} - \mathcal{A}\| \leq \epsilon$, while
 558 $QE(\mathcal{A}; \mathcal{K})$ can be on the order of $\ln(R)$ for all $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$. Indeed, this follows directly from a
 559 result in [25], which we restate as Lemma 6 for convenience.

560 **Proposition 2.** Let $f : \mathbb{R}^2 \rightarrow \mathbb{R}_{\geq 0}$ be monotonically increasing in its second variable. Suppose
 561 that the following statement holds: if a tensor $\mathcal{A} \in \mathbb{R}^{D_1 \times \dots \times D_N}$ satisfies $QE(\mathcal{A}; \mathcal{K}) \leq f(\|\mathcal{A}\|, \epsilon)$
 562 for all canonical partitions $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$ (Definition 2), then there exists an assignment for the
 563 tensors constituting the locally connected tensor network (defined in Section 3.1) for which $\mathcal{W}_{\text{TN}} \in$
 564 $\mathbb{R}^{D_1 \times \dots \times D_N}$ upholds:

$$\|\mathcal{W}_{\text{TN}} - \mathcal{A}\| \leq \epsilon.$$

565 Then, for any $\mathcal{A} \in \mathbb{R}^{D_1 \times \dots \times D_N}$, as the desired approximation error ϵ goes to zero, so does the
 566 sufficient condition on entanglements, i.e.:

$$\lim_{\epsilon \rightarrow 0} f(\|\mathcal{A}\|, \epsilon) = 0.$$

567 *Proof.* Suppose otherwise, i.e. that there exists some $a > 0$ such that $\lim_{\epsilon \rightarrow 0} f(a, \epsilon) =$
 568 $\inf_{\epsilon > 0} f(a, \epsilon) = c > 0$. Let \mathcal{A} be a tensor with $\|\mathcal{A}\| = a$ such that $QE(\mathcal{A}; \mathcal{K}) < c$ for all
 569 canonical partitions $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$. Then by assumption, for all $\epsilon > 0$ there exist a tensor generated
 570 by the locally connected tensor network, which we denote by $\mathcal{W}_{\text{TN}}(\epsilon) \in \mathbb{R}^{D_1 \times \dots \times D_N}$, that satisfies:

$$\|\mathcal{W}_{\text{TN}}(\epsilon) - \mathcal{A}\| \leq \epsilon.$$

571 By Lemma 4, for all $\epsilon > 0$ we have that

$$\text{rank}(\llbracket \mathcal{W}_{\text{TN}}(\epsilon); \mathcal{K} \rrbracket) \leq R,$$

572 so by the lower semicontinuity of the matrix rank, we have that $\text{rank}(\llbracket \mathcal{A}; \mathcal{K} \rrbracket) \leq R$ as well. But, there
 573 are tensors for which this leads to a contradiction, namely tensors with arbitrary low entanglement
 574 across all partitions but nearly maximal rank. Indeed, consider tensors of the form

$$\mathcal{Q}(\delta) = a \cdot \frac{\otimes_{n=1}^N \mathbf{v}^{(n)} + \delta \cdot \mathcal{B}}{\|\otimes_{n=1}^N \mathbf{v}^{(n)} + \delta \cdot \mathcal{B}\|} \in \mathbb{R}^{D_1 \times \dots \times D_N},$$

575 where $\{\mathbf{v}^{(n)} \in \mathbb{R}^{D_n}\}_{n=1}^N$ are (non-zero) vectors, $\delta > 0$ and $\mathcal{B} \in \mathbb{R}^{D_1 \times \dots \times D_N}$ is some tensor with
576 matricizations of maximal rank across all canonical partitions (for a proof of the existence of such a
577 tensor see Claim 3 in [35]). Note that $\|\mathcal{Q}(\delta)\| = a$ for all $\delta > 0$. By the triangle inequality have

$$\left\| \frac{\mathcal{Q}(\delta)}{a} - \frac{\otimes_{n=1}^N \mathbf{v}^{(n)}}{\|\otimes_{n=1}^N \mathbf{v}^{(n)}\|} \right\| \leq \frac{\delta \cdot \|\mathcal{B}\|}{\|\otimes_{n=1}^N \mathbf{v}^{(n)} + \delta \cdot \mathcal{B}\|} + \left\| \frac{\otimes_{n=1}^N \mathbf{v}^{(n)}}{\|\otimes_{n=1}^N \mathbf{v}^{(n)} + \delta \cdot \mathcal{B}\|} - \frac{\otimes_{n=1}^N \mathbf{v}^{(n)}}{\|\otimes_{n=1}^N \mathbf{v}^{(n)}\|} \right\|,$$

578 and so

$$\lim_{\delta \rightarrow 0} \left\| \frac{\mathcal{Q}(\delta)}{a} - \frac{\otimes_{n=1}^N \mathbf{v}^{(n)}}{\|\otimes_{n=1}^N \mathbf{v}^{(n)}\|} \right\| = 0.$$

579 Thus, from Lemma 8 we know that

$$\lim_{\delta \rightarrow 0} \left| QE\left(\frac{\mathcal{Q}(\delta)}{a}; \mathcal{K}\right) - QE\left(\frac{\otimes_{n=1}^N \mathbf{v}^{(n)}}{\|\otimes_{n=1}^N \mathbf{v}^{(n)}\|}; \mathcal{K}\right) \right| = 0$$

580 for all $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$. Furthermore, for any $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$:

$$QE\left(\frac{\otimes_{n=1}^N \mathbf{v}^{(n)}}{\|\otimes_{n=1}^N \mathbf{v}^{(n)}\|}; \mathcal{K}\right) = 0,$$

581 and therefore by Theorem 1 for sufficiently small $\delta > 0$ we have

$$QE(\mathcal{Q}(\delta); \mathcal{K}) < c,$$

582 for all $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$. However, $\text{rank}(\llbracket \mathcal{Q}(\delta); \mathcal{K} \rrbracket)$ is (nearly) maximal for all canonical partitions.
583 Indeed, by the triangle inequality for matrix rank we get

$$\text{rank}(\llbracket \mathcal{Q}(\delta); \mathcal{K} \rrbracket) \geq \text{rank}(\llbracket \mathcal{B}; \mathcal{K} \rrbracket) - \text{rank}\left(\llbracket \otimes_{n=1}^N \mathbf{v}^{(n)}; \mathcal{K} \rrbracket\right) = D_{\mathcal{K}} - 1,$$

584 where $D_{\mathcal{K}} := \min\{\prod_{n \in \mathcal{K}} D_n, \prod_{n \in \mathcal{K}^c} D_n\}$. □

585 E Efficiently Computing Entanglements of the Empirical Data Tensor

586 For a given tensor, its entanglement with respect to a partition of axes (Definition 1) is determined by
587 the singular values of its arrangement as a matrix according to the partition. Since the empirical data
588 tensor \mathcal{D}_{emp} (Equation (2)) has size exponential in the number of features N , it is infeasible to naively
589 compute its entanglement (or even explicitly store it in memory). Fortunately, as shown in [40], the
590 specific form of the empirical data tensor admits an efficient algorithm for computing entanglements,
591 without explicitly manipulating an exponentially large matrix. Specifically, the algorithm runs in
592 $\mathcal{O}(DNM^2 + M^3)$ time and requires $\mathcal{O}(DNM + M^2)$ memory, where $D := \max_{n \in [N]} D_n$ is
593 the maximal feature dimension (*i.e.* axis length of \mathcal{D}_{emp}), N is the number of features in the data
594 (*i.e.* number of axes that \mathcal{D}_{emp} has), and M is the number of training instances. For completeness,
595 we outline the method in Algorithm 1 while referring the interested reader to Appendix A in [40] for
596 further details.

597 F Practical Algorithm via Surrogate for Entanglement

598 To efficiently implement the scheme from Section 5.1, we replace entanglement with a surrogate
599 measure of dependence. The surrogate is based on the Pearson correlation coefficient for multivariate
600 features [45],⁶ and its agreement with entanglement is demonstrated empirically in Appendix H.
601 Theoretically supporting this agreement is left for future work.

602 **Definition 4.** Given a set of $M \in \mathbb{N}$ instances $\mathcal{X} := \{\mathbf{x}^{(1,m)}, \dots, \mathbf{x}^{(N,m)}\} \in (\mathbb{R}^D)^N$, denote
603 by $p_{n,n'}$ the multivariate Pearson correlation between features $n, n' \in [N]$. For $\mathcal{K} \subseteq [N]$, the
604 *surrogate entanglement* of \mathcal{X} with respect to the partition $(\mathcal{K}, \mathcal{K}^c)$, denoted $SE(\mathcal{X}; \mathcal{K})$, is the sum of
605 absolute values of Pearson correlation coefficients between pairs of features, the first belonging to \mathcal{K}
606 and the second to $\mathcal{K}^c := [N] \setminus \mathcal{K}$. That is, $SE(\mathcal{X}; \mathcal{K}) := \sum_{n \in \mathcal{K}, n' \in \mathcal{K}^c} |p_{n,n'}|$.

607 As shown by Proposition 3 below, replacing entanglement with surrogate entanglement in the scheme
608 from Section 5.1 converts each search for a balanced partition minimizing average entanglement into
609 a *minimum balanced cut problem*. Although the minimum balanced cut problem is NP-hard (see,
610 *e.g.*, [23]), it enjoys a wide array of well-established approximation tools, particularly ones designed
611 for large scale [29, 56]. We therefore obtain a practical algorithm for enhancing the suitability of a
612 data distribution to locally connected neural networks — see Algorithm 2.

⁶For completeness, Appendix G provides a formal definition of the multivariate Pearson correlation.

Algorithm 1 Entanglement Computation for the Empirical Data Tensor

1: **Input:** $\mathcal{X} := \{\mathbf{x}^{(1,m)}, \dots, \mathbf{x}^{(N,m)}\}_{m=1}^M$ — $M \in \mathbb{N}$ data instances comprising $N \in \mathbb{N}$ features each, $\mathcal{K} \subseteq [N]$ — subset of feature indices

2: **Output:** $QE(\mathcal{D}_{\text{emp}}; \mathcal{K})$

3: Compute $\mathbf{G}^{(\mathcal{K})}, \mathbf{G}^{(\mathcal{K}^c)} \in \mathbb{R}^{M \times M}$ given element-wise by:

$$\forall i, j \in [M] : \mathbf{G}_{i,j}^{(\mathcal{K})} = y^{(i)} y^{(j)} \cdot \langle \otimes_{n \in \mathcal{K}} \mathbf{x}^{(n,i)}, \otimes_{n \in \mathcal{K}} \mathbf{x}^{(n,j)} \rangle = y^{(i)} y^{(j)} \cdot \prod_{n \in \mathcal{K}} \langle \mathbf{x}^{(n,i)}, \mathbf{x}^{(n,j)} \rangle$$

$$\forall i, j \in [M] : \mathbf{G}_{i,j}^{(\mathcal{K}^c)} = \langle \otimes_{n \in \mathcal{K}^c} \mathbf{x}^{(n,i)}, \otimes_{n \in \mathcal{K}^c} \mathbf{x}^{(n,j)} \rangle = \prod_{n \in \mathcal{K}^c} \langle \mathbf{x}^{(n,i)}, \mathbf{x}^{(n,j)} \rangle$$

4: Compute eigenvalue decompositions of $\mathbf{G}^{(\mathcal{K})}$ and $\mathbf{G}^{(\mathcal{K}^c)}$, i.e.:

$$\mathbf{G}^{(\mathcal{K})} = \mathbf{U}^{(\mathcal{K})} \mathbf{S}^{(\mathcal{K})} (\mathbf{U}^{(\mathcal{K})})^\top$$

$$\mathbf{G}^{(\mathcal{K}^c)} = \mathbf{U}^{(\mathcal{K}^c)} \mathbf{S}^{(\mathcal{K}^c)} (\mathbf{U}^{(\mathcal{K}^c)})^\top$$

where $\mathbf{U}^{(\mathcal{K})}, \mathbf{U}^{(\mathcal{K}^c)} \in \mathbb{R}^{M \times M}$ are orthogonal matrices and $\mathbf{S}^{(\mathcal{K})}, \mathbf{S}^{(\mathcal{K}^c)} \in \mathbb{R}^{M \times M}$ are diagonal holding the eigenvalues of $\mathbf{G}^{(\mathcal{K})}$ and $\mathbf{G}^{(\mathcal{K}^c)}$, respectively

5: Compute $\mathbf{Q} = (\mathbf{S}^{(\mathcal{K})})^{\frac{1}{2}} (\mathbf{U}^{(\mathcal{K})})^\top \mathbf{U}^{(\mathcal{K}^c)} (\mathbf{S}^{(\mathcal{K}^c)})^{\frac{1}{2}} \in \mathbb{R}^{M \times M}$

6: Compute a singular value decomposition of \mathbf{Q} to obtain its singular values $\sigma_1(\mathbf{Q}), \dots, \sigma_M(\mathbf{Q})$

7: Let $\rho_m := \sigma_m^2(\mathbf{Q}) / \sum_{m'=1}^M \sigma_{m'}^2(\mathbf{Q})$ for $m \in [M]$

8: **return** $QE(\mathcal{D}_{\text{emp}}; \mathcal{K}) = -\sum_{m=1}^M \rho_m \ln(\rho_m)$ (if $\mathbf{Q} = 0$, then return 0)

613 **Proposition 3.** For any $\bar{\mathcal{K}} \subseteq [N]$ of even size, the following optimization problem can be framed as
 614 a minimum balanced cut problem over a complete graph with $|\bar{\mathcal{K}}|$ vertices:

$$\min_{\mathcal{K} \subset \bar{\mathcal{K}}, |\mathcal{K}| = |\bar{\mathcal{K}}|/2} \frac{1}{2} [SE(\mathcal{X}; \mathcal{K}) + SE(\mathcal{X}; \bar{\mathcal{K}} \setminus \mathcal{K})]. \quad (4)$$

615 Specifically, there exists a complete undirected weighted graph with vertices $\bar{\mathcal{K}}$ and edge weights
 616 $w : \bar{\mathcal{K}} \times \bar{\mathcal{K}} \rightarrow \mathbb{R}$ such that for any $\mathcal{K} \subset \bar{\mathcal{K}}$, the weight of the cut in the graph induced by \mathcal{K} —
 617 $\sum_{n \in \mathcal{K}, n' \in \bar{\mathcal{K}} \setminus \mathcal{K}} w(\{n, n'\})$ — is equal, up to an additive constant, to the term minimized in Equa-
 618 tion (4), i.e. to $\frac{1}{2} [SE(\mathcal{X}; \mathcal{K}) + SE(\mathcal{X}; \bar{\mathcal{K}} \setminus \mathcal{K})]$.

619 *Proof.* Consider the complete undirected graph whose vertices are $\bar{\mathcal{K}}$ and where the weight of an
 620 edge $\{n, n'\} \in \bar{\mathcal{K}} \times \bar{\mathcal{K}}$ is $w(\{n, n'\}) = |p_{n,n'}|$ (recall that $p_{n,n'}$ stands for the multivariate Pearson
 621 correlation between features n and n' in \mathcal{X}). For any $\mathcal{K} \subset \bar{\mathcal{K}}$ it holds that:

$$\sum_{n \in \mathcal{K}, n' \in \bar{\mathcal{K}} \setminus \mathcal{K}} w(\{n, n'\}) = \frac{1}{2} [SE(\mathcal{X}; \mathcal{K}) + SE(\mathcal{X}; \bar{\mathcal{K}} \setminus \mathcal{K})] - \frac{1}{2} SE(\mathcal{X}; \bar{\mathcal{K}}),$$

622 where $\frac{1}{2} SE(\mathcal{X}; \bar{\mathcal{K}})$ does not depend on \mathcal{K} . This concludes the proof. \square

623 G Definition of the Multivariate Pearson Correlation from [45]

624 Appendix F introduces a surrogate measure for entanglement based on the multivariate Pearson
 625 correlation from [45]. For completeness, this appendix provides its formal definition.

626 Given a set of $M \in \mathbb{N}$ instances $\mathcal{X} := \{(\mathbf{x}^{(1,m)}, \dots, \mathbf{x}^{(N,m)}) \in (\mathbb{R}^D)^N\}_{m=1}^M$, let $\Sigma^{(n)}$ be the
 627 empirical covariance matrix of feature $n \in [N]$ and $\Sigma^{(n,n')}$ be the empirical cross-covariance matrix

Algorithm 2 Enhancing Suitability of Data to Locally Connected Neural Networks

- 1: **Input:** $\mathcal{X} := \{(\mathbf{x}^{(1,m)}, \dots, \mathbf{x}^{(N,m)})\}_{m=1}^M$ — $M \in \mathbb{N}$ data instances comprising $N \in \mathbb{N}$ features
 - 2: **Output:** Permutation $\pi : [N] \rightarrow [N]$ to apply to feature indices
-
- 3: Let $\mathcal{K}_{0,1} := [N]$ and denote $L := \log_2(N)$
 - 4: # We assume for simplicity that N is a power of two, otherwise one may add constant features
 - 5: **for** $l = 0, \dots, L - 1$, $n = 1, \dots, 2^l$ **do**
 - 6: Using a reduction to a minimum balanced cut problem (Proposition 3), find an approximate solution $\mathcal{K}_{l+1,2n-1} \subset \mathcal{K}_{l,n}$ for:

$$\min_{\mathcal{K} \subset \mathcal{K}_{l,n}, |\mathcal{K}|=|\mathcal{K}_{l,n}|/2} \frac{1}{2} [SE(\mathcal{X}; \mathcal{K}) + SE(\mathcal{X}; \mathcal{K}_{l,n} \setminus \mathcal{K})]$$
 - 7: Let $\mathcal{K}_{l+1,2n} := \mathcal{K}_{l,n} \setminus \mathcal{K}_{l+1,2n-1}$
 - 8: **end for**
 - 9: # At this point, $\mathcal{K}_{L,1}, \dots, \mathcal{K}_{L,N}$ each contain a single feature index
 - 10: **return** π that maps $k \in \mathcal{K}_{L,n}$ to n , for every $n \in [N]$
-

628 of features $n, n' \in [N]$, i.e.:

$$\Sigma^{(n)} := \frac{1}{M} \sum_{m=1}^M (\mathbf{x}^{(n,m)} - \mu^{(n)}) \otimes (\mathbf{x}^{(n,m)} - \mu^{(n)}),$$

$$\Sigma^{(n,n')} := \frac{1}{M} \sum_{m=1}^M (\mathbf{x}^{(n,m)} - \mu^{(n)}) \otimes (\mathbf{x}^{(n',m)} - \mu^{(n')}),$$

629 where $\mu^{(n)} := \frac{1}{M} \sum_{m=1}^M \mathbf{x}^{(n,m)}$ for $n \in [N]$. With this notation, the multivariate Pearson correlation
 630 of $n, n' \in [N]$ from [45] is defined by $p_{n,n'} := \text{trace}(\Sigma^{(n,n')}) / \text{trace}((\Sigma^{(n)} \Sigma^{(n')})^{1/2})$.

631 H Entanglement and Surrogate Entanglement Are Strongly Correlated

632 In Appendix F, we introduced a surrogate entanglement measure (Definition 4) to facilitate efficient
 633 implementation of the feature arrangement search scheme from Section 5.1. Figure 5 supports
 634 the viability of the chosen surrogate, demonstrating empirically that it is strongly correlated with
 635 entanglement of the empirical data tensor (Definition 1 and Equation (2)).

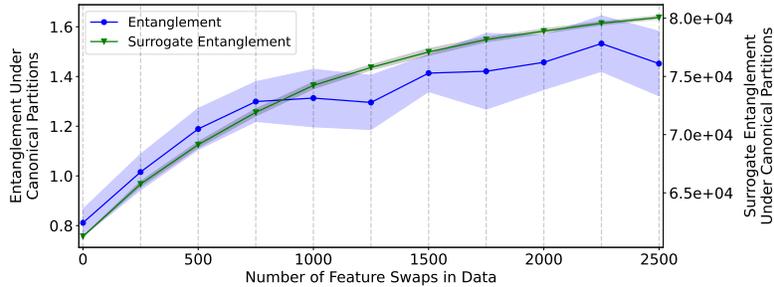


Figure 5: Surrogate entanglement (Definition 4) is strongly correlated with the entanglement (Definition 1) of the empirical data tensor. Presented are average entanglement and average surrogate entanglement under canonical partitions, admitted by the Speech Commands audio datasets [62] considered in Figure 3. Remarkably, the Pearson correlation between the quantities is 0.977. For further details see caption of Figure 3 as well as Appendix J.2.

636 I Extension to Arbitrary Dimensional Models and Data

637 In this appendix, we extend our theoretical analysis and experiments, including the algorithm for
 638 enhancing the suitability of data to locally connected neural networks, from one-dimensional (se-
 639 quential) models and data to P -dimensional models and data (such as two-dimensional image data or

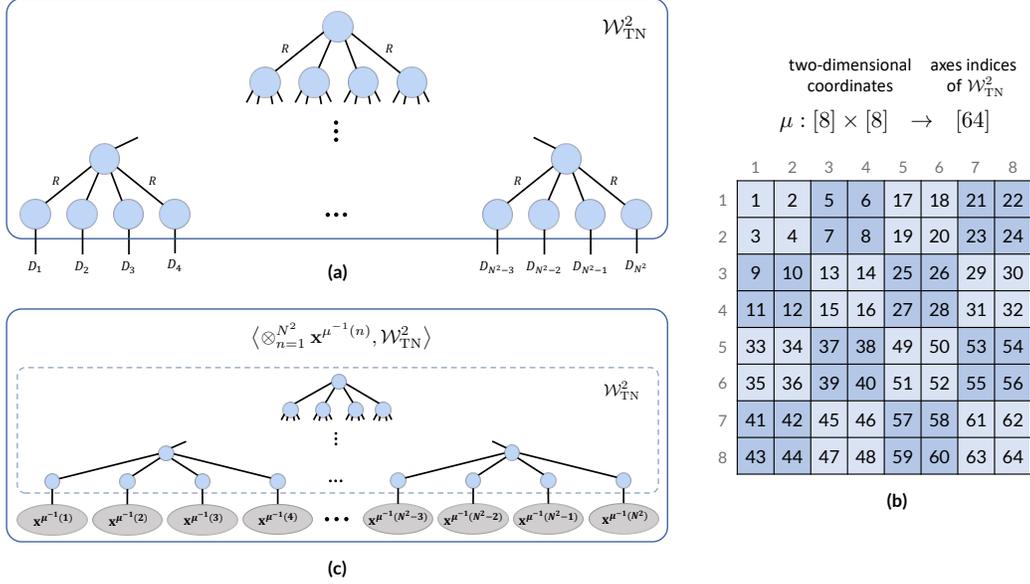


Figure 6: The analyzed tensor network equivalent to a locally connected neural network operating over P -dimensional data, for $P = 2$. (a) The tensor network adheres to a perfect 2^P -ary tree connectivity with N^P leaf nodes, where $N = 2^L$ for some $L \in \mathbb{N}$, and generates $\mathcal{W}_{\text{TN}}^P \in \mathbb{R}^{D_1 \times \dots \times D_{N^P}}$. Axes corresponding to open edges are indexed such that open edges descendant to any node of the tree have contiguous indices. The lengths of axes corresponding to inner (non-open) edges are equal to $R \in \mathbb{N}$, referred to as the width of the tensor network. (b) Exemplar $\mu : [N]^P \rightarrow [N^P]$ compatible with the locally connected tensor network (Definition 5), mapping P -dimensional coordinates to axes indices of $\mathcal{W}_{\text{TN}}^P$. (c) Contracting $\mathcal{W}_{\text{TN}}^P$ with vectors $\{\mathbf{x}^{(n_1, \dots, n_P)}\}_{n_1, \dots, n_P \in [N]}$ according to a compatible μ produces $\langle \otimes_{n=1}^{N^P} \mathbf{x}^{\mu^{-1}(n)}, \mathcal{W}_{\text{TN}}^P \rangle$. Performing these contractions can be viewed as a forward pass of a certain locally connected neural network (with polynomial non-linearity) over the data instance $\{\mathbf{x}^{(n_1, \dots, n_P)}\}_{n_1, \dots, n_P \in [N]}$ (see, e.g., [11, 9, 35, 49]).

640 three-dimensional video data), for $P \in \mathbb{N}$. Specifically, Appendix I.1 extends Section 3, Appendix I.2
 641 extends Section 4 and Appendix I.3 extends Section 5.

642 To ease presentation, we consider P -dimensional data instances whose feature vectors are associated
 643 with coordinates $(n_1, \dots, n_P) \in [N]^P$, where $N = 2^L$ for some $L \in \mathbb{N}$ (if this is not the case we
 644 may add constant features as needed).

645 I.1 Low Entanglement Under Canonical Partitions Is Necessary and Sufficient for Fitting 646 Tensor

647 We introduce the locally connected tensor network equivalent to a locally connected neural network
 648 that operates over P -dimensional data (Appendix I.1.1). Subsequently, we establish a necessary
 649 and sufficient condition required for it to fit a given tensor (Appendix I.1.2), generalizing the results
 650 of Section 3.2.

651 I.1.1 Tensor Network Equivalent to a Locally Connected Neural Network

652 For P -dimensional data, the locally connected tensor network we consider (defined in Section 3.1
 653 for one-dimensional data) has an underlying perfect 2^P -ary tree graph of height L . We denote the
 654 tensor it generates by $\mathcal{W}_{\text{TN}}^P \in \mathbb{R}^{D_1 \times \dots \times D_{N^P}}$. Figure 6(a) provides its diagrammatic definition. As in
 655 the one-dimensional case, the lengths of axes corresponding to inner edges are taken to be $R \in \mathbb{N}$,
 656 referred to as the width of the tensor network.

657 The axes of $\mathcal{W}_{\text{TN}}^P$ are associated with P -dimensional coordinates through a bijective function $\mu :$
 658 $[N]^P \rightarrow [N^P]$.

659 **Definition 5.** We say that a bijective function $\mu : [N]^P \rightarrow [N^P]$ is *compatible* with the locally
 660 connected tensor network if, for any node in the tensor network, the coordinates mapped to indices
 661 of $\mathcal{W}_{\text{TN}}^P$'s axes descendant to that node form a contiguous P -dimensional cubic block in $[N]^P$ (e.g.,

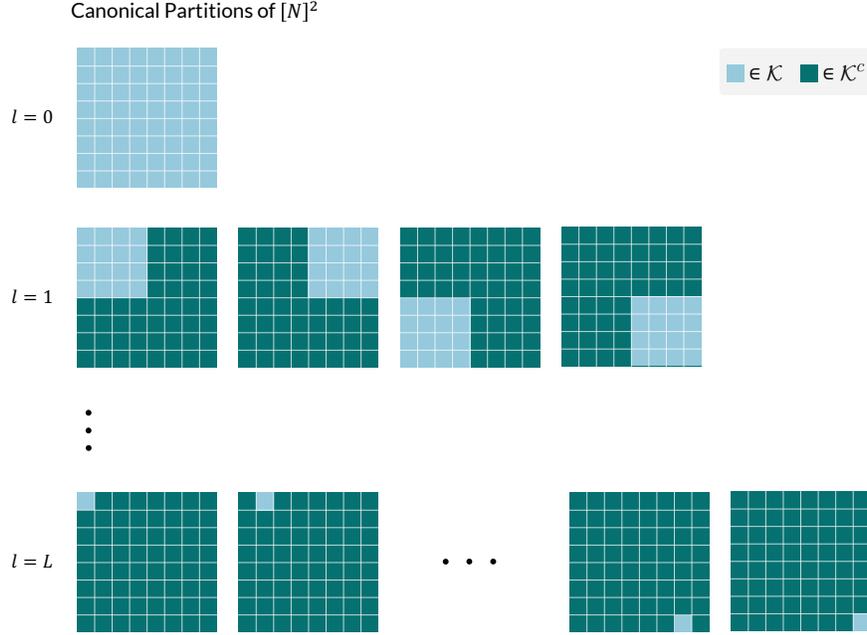


Figure 7: The canonical partitions of $[N]^P$, for $P = 2$ and $N = 2^L$ with $L \in \mathbb{N}$. Every $l \in \{0, \dots, L\}$ contributes $2^{l \cdot P}$ canonical partitions, each induced by $\mathcal{K} = \times_{p=1}^P \{2^{L-l} \cdot (n_p - 1) + 1, \dots, 2^{L-l} \cdot n_p\}$ for $n_1, \dots, n_P \in [2^l]$.

662 square block when $P = 2$) — see Figure 6(b) for an illustration. With slight abuse of notation, for
 663 $\mathcal{K} \subseteq [N]^P$ we denote $\mu(\mathcal{K}) := \{\mu(n_1, \dots, n_P) : (n_1, \dots, n_P) \in \mathcal{K}\} \subseteq [N]^P$.

664 Contracting the locally connected tensor network with $\{\mathbf{x}^{(n_1, \dots, n_P)} \in \mathbb{R}^{D_{\mu(n_1, \dots, n_P)}}\}_{n_1, \dots, n_P \in [N]}$
 665 according to a compatible μ , as depicted in Figure 6(c), can be viewed as a forward pass of the data
 666 instance $\{\mathbf{x}^{(n_1, \dots, n_P)}\}_{n_1, \dots, n_P \in [N]}$ through a locally connected neural network (with polynomial
 667 non-linearity), which produces the scalar $\langle \otimes_{n=1}^{N^P} \mathbf{x}^{\mu^{-1}(n)}, \mathcal{W}_{\text{TN}}^P \rangle$ (see, *e.g.*, [11, 9, 35, 49]).

668 I.1.2 Necessary and Sufficient Condition for Fitting Tensor

669 The ability of the locally connected tensor network, defined in Appendix I.1.1, to fit (*i.e.* represent) a
 670 tensor is determined by the entanglements that the tensor admits under partitions of its axes, induced
 671 by the following canonical partitions of $[N]^P$.

672 **Definition 6.** The *canonical partitions* of $[N]^P$, illustrated in Figure 7 for $P = 2$, are:⁷

$$\mathcal{C}_N^P := \left\{ (\mathcal{K}, \mathcal{K}^c) : \mathcal{K} = \times_{p=1}^P \{2^{L-l} \cdot (n_p - 1) + 1, \dots, 2^{L-l} \cdot n_p\}, \right. \\ \left. l \in \{0, \dots, L\}, n_1, \dots, n_P \in [2^l] \right\}.$$

673 With the definition of canonical partitions for P -dimensional data in place, Theorem 3 generalizes The-
 674 orem 1. In particular, suppose that $\mathcal{W}_{\text{TN}}^P$ — the tensor generated by the locally connected tensor
 675 network — well-approximates $\mathcal{A} \in \mathbb{R}^{D_1 \times \dots \times D_{N^P}}$. Then, given a compatible $\mu : [N]^P \rightarrow [N]^P$
 676 (Definition 5), Theorem 3 establishes that the entanglement of \mathcal{A} with respect to $\mu(\mathcal{K})$, where \mathcal{K} is a
 677 canonical partition, cannot be much larger than $\ln(R)$, whereas the entanglement attainable by an
 678 arbitrary tensor can be linear in N^P .

679 In the other direction, Theorem 4 implies that low entanglement under partitions of axes induced
 680 by canonical partitions of $[N]^P$ is not only necessary for a tensor to be fit by the locally connected
 681 tensor network, but also sufficient.

⁷For sets $\mathcal{S}_1, \dots, \mathcal{S}_P$, we denote their Cartesian product by $\times_{p=1}^P \mathcal{S}_p$.

682 **Theorem 3.** Let $\mathcal{W}_{TN}^P \in \mathbb{R}^{D_1 \times \dots \times D_{N^P}}$ be a tensor generated by the locally connected tensor
683 network defined in Appendix I.1.1, and $\mu : [N]^P \rightarrow [N^P]$ be a compatible map from P -dimensional
684 coordinates to axes indices of \mathcal{W}_{TN}^P (Definition 5). For any $\mathcal{A} \in \mathbb{R}^{D_1 \times \dots \times D_{N^P}}$ and $\epsilon \in [0, \|\mathcal{A}\|/4]$, if
685 $\|\mathcal{W}_{TN}^P - \mathcal{A}\| \leq \epsilon$, then for all canonical partitions $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N^P$ (Definition 6):

$$QE(\mathcal{A}; \mu(\mathcal{K})) \leq \ln(R) + \frac{2\epsilon}{\|\mathcal{A}\|} \cdot \ln(D_{\mu(\mathcal{K})}) + 2\sqrt{\frac{2\epsilon}{\|\mathcal{A}\|}}, \quad (5)$$

686 where $D_{\mu(\mathcal{K})} := \min\{\prod_{n \in \mu(\mathcal{K})} D_n, \prod_{n \in \mu(\mathcal{K})^c} D_n\}$. In contrast, there exists $\mathcal{A}' \in \mathbb{R}^{D_1 \times \dots \times D_{N^P}}$
687 such that for all canonical partitions $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N^P$:

$$QE(\mathcal{A}'; \mu(\mathcal{K})) \geq \min\{|\mathcal{K}|, |\mathcal{K}^c|\} \cdot \ln(\min_{n \in [N^P]} D_n). \quad (6)$$

688 *Proof sketch (proof in Appendix K.6).* The proof is analogous to that of Theorem 1. \square

689 **Theorem 4.** Let $\mathcal{A} \in \mathbb{R}^{D_1 \times \dots \times D_{N^P}}$ and $\epsilon > 0$. Suppose that for all canonical partitions $(\mathcal{K}, \mathcal{K}^c) \in$
690 \mathcal{C}_N^P (Definition 6) it holds that $QE(\mathcal{A}; \mu(\mathcal{K})) \leq \frac{\epsilon^2}{(2N^P-3)\|\mathcal{A}\|^2} \cdot \ln(R)$, where $\mu : [N]^P \rightarrow [N^P]$ is
691 compatible with the locally connected tensor network (Definition 5). Then, there exists an assignment
692 for the tensors constituting the locally connected tensor network (defined in Appendix I.1.1) such that
693 it generates $\mathcal{W}_{TN}^P \in \mathbb{R}^{D_1 \times \dots \times D_{N^P}}$ satisfying:

$$\|\mathcal{W}_{TN}^P - \mathcal{A}\| \leq \epsilon.$$

694 *Proof sketch (proof in Appendix K.7).* The claim follows through a reduction from the locally con-
695 nected tensor network for P -dimensional data to that for one-dimensional data (defined in Section 3.1),
696 i.e. from perfect 2^P -ary to perfect binary tree tensor networks. Specifically, we consider a modified
697 locally connected tensor network for one-dimensional data, where axes corresponding to different
698 inner edges can vary in length (as opposed to all having length R). We then show, by arguments anal-
699 ogous to those in the proof of Theorem 2, that it can approximate \mathcal{A} while having certain inner axes,
700 related to the canonical partitions of $[N]^P$, of lengths at most R . The proof concludes by establishing
701 that, any tensor represented by such a locally connected tensor network for one-dimensional data can
702 be represented via the locally connected tensor network for P -dimensional data (where the length of
703 each axis corresponding to an inner edge is R). \square

704 I.2 Low Entanglement Under Canonical Partitions Is Necessary and Sufficient for Accurate 705 Prediction

706 In this appendix, we consider the locally connected tensor network from Appendix I.1.1 in a machine
707 learning setting, and extend the results and experiments of Section 4 from one-dimensional to
708 P -dimensional models and data.

709 I.2.1 Accurate Prediction Is Equivalent to Fitting Data Tensor

710 The locally connected tensor network generating $\mathcal{W}_{TN}^P \in \mathbb{R}^{D_1 \times \dots \times D_{N^P}}$ is equivalent to a locally
711 connected neural network operating over P -dimensional data (see Appendix I.1.1). Specifically, a for-
712 ward pass of the latter over $\{\mathbf{x}^{(n_1, \dots, n_P)} \in \mathbb{R}^{D_{\mu(n_1, \dots, n_P)}}\}_{n_1, \dots, n_P \in [N]}$ yields $\langle \otimes_{n=1}^{N^P} \mathbf{x}^{\mu^{-1}(n)}, \mathcal{W}_{TN}^P \rangle$,
713 for a compatible $\mu : [N]^P \rightarrow [N^P]$ (Definition 5). Suppose we are given a training set of labeled
714 instances $\{\{\mathbf{x}^{((n_1, \dots, n_P), m)}\}_{n_1, \dots, n_P \in [N]}, y^{(m)}\}_{m=1}^M$ drawn i.i.d. from some distribution, where
715 $y^{(m)} \in \{1, -1\}$ for $m \in [M]$. Learning the parameters of the neural network through the soft-margin
716 support vector machine (SVM) objective amounts to optimizing:

$$\min_{\|\mathcal{W}_{TN}^P\| \leq B} \frac{1}{M} \sum_{m=1}^M \max\left\{0, 1 - y^{(m)} \langle \otimes_{n=1}^{N^P} \mathbf{x}^{(\mu^{-1}(n), m)}, \mathcal{W}_{TN}^P \rangle\right\},$$

717 for a predetermined constant $B > 0$. This objective generalizes Equation (1) from one-dimensional to
718 P -dimensional model and data. Assume that instances are normalized, i.e. $\|\mathbf{x}^{((n_1, \dots, n_P), m)}\| \leq 1$ for
719 all $n_1, \dots, n_P \in [N]$, $m \in [M]$, and that $B \leq 1$. By a derivation analogous to that succeeding Equa-
720 tion (1) in Section 4.1, it follows that minimizing the SVM objective is equivalent to minimizing

721 $\left\| \frac{\mathcal{W}_{TN}^P}{\|\mathcal{W}_{TN}^P\|} - \frac{\mathcal{D}_{\text{emp}}^P}{\|\mathcal{D}_{\text{emp}}^P\|} \right\|$, where:

$$\mathcal{D}_{\text{emp}}^P := \frac{1}{M} \sum_{m=1}^M y^{(m)} \cdot \otimes_{n=1}^{N^P} \mathbf{x}^{(\mu^{-1}(n), m)} \quad (7)$$

722 extends the notion of *empirical data tensor* to P -dimensional data. In other words, the accuracy
 723 achievable over the training data is determined by the extent to which $\frac{\mathcal{W}_{\text{TN}}^P}{\|\mathcal{W}_{\text{TN}}^P\|}$ can fit the normalized
 724 empirical data tensor $\frac{\mathcal{D}_{\text{emp}}^P}{\|\mathcal{D}_{\text{emp}}^P\|}$.

725 The same arguments apply to the population loss, in which case $\mathcal{D}_{\text{emp}}^P$ is replaced by the *population*
 726 *data tensor*:

$$\mathcal{D}_{\text{pop}}^P := \mathbb{E}_{\{\mathbf{x}^{(n_1, \dots, n_P)}\}_{n_1, \dots, n_P \in [N]; y}} \left[y \cdot \otimes_{n=1}^{N^P} \mathbf{x}^{\mu^{-1}(n)} \right]. \quad (8)$$

727 The achievable accuracy over the population is therefore determined by the extent to which $\frac{\mathcal{W}_{\text{TN}}^P}{\|\mathcal{W}_{\text{TN}}^P\|}$
 728 can fit the normalized population data tensor $\frac{\mathcal{D}_{\text{pop}}^P}{\|\mathcal{D}_{\text{pop}}^P\|}$. Accordingly, we refer to the minimal distance
 729 from it as the *suboptimality in achievable accuracy*, generalizing Definition 3 from Section 4.1.

730 **Definition 7.** In the context of the classification setting above, the *suboptimality in achievable*
 731 *accuracy* is:

$$\text{SubOpt}^P := \min_{\mathcal{W}_{\text{TN}}^P} \left\| \frac{\mathcal{W}_{\text{TN}}^P}{\|\mathcal{W}_{\text{TN}}^P\|} - \frac{\mathcal{D}_{\text{pop}}^P}{\|\mathcal{D}_{\text{pop}}^P\|} \right\|.$$

732 I.2.2 Necessary and Sufficient Condition for Accurate Prediction

733 In the classification setting of Appendix I.2.1, by invoking Theorems 3 and 4 from Appendix I.1.2,
 734 we conclude that the suboptimality in achievable accuracy is small if and only if the population
 735 (empirical) data tensor $\mathcal{D}_{\text{pop}}^P$ ($\mathcal{D}_{\text{emp}}^P$) admits low entanglement under the canonical partitions of
 736 features (Definition 6). Specifically, we establish Corollary 3, Proposition 4 and Corollary 4, which
 737 generalize Corollary 1, Proposition 4 and Corollary 2 from Section 4.2, respectively, to P -dimensional
 738 model and data.

739 **Corollary 3.** Consider the classification setting of Appendix I.2.1, and let $\epsilon \in [0, 1/4]$. If there exists
 740 a canonical partition $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N^P$ (Definition 6) under which $QE(\mathcal{D}_{\text{pop}}^P; \mu(\mathcal{K})) > \ln(R) + 2\epsilon \cdot$
 741 $\ln(D_{\mu(\mathcal{K})}) + 2\sqrt{2}\epsilon$, where $D_{\mu(\mathcal{K})} := \min\{\prod_{n \in \mu(\mathcal{K})} D_n, \prod_{n \in \mu(\mathcal{K})^c} D_n\}$, then:

$$\text{SubOpt}^P > \epsilon.$$

742 Conversely, if for all $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N^P$ it holds that $QE(\mathcal{D}_{\text{pop}}^P; \mu(\mathcal{K})) \leq \frac{\epsilon^2}{8N^P - 12} \cdot \ln(R)$, then:

$$\text{SubOpt}^P \leq \epsilon.$$

743 *Proof.* Implied by Theorems 3 and 4 after accounting for $\mathcal{W}_{\text{TN}}^P$ being normalized in the suboptimality
 744 in achievable accuracy, as done in the proof of Corollary 1. \square

745 **Proposition 4.** Consider the classification setting of Appendix I.2.1, and let $\delta \in (0, 1)$ and
 746 $\gamma > 0$. If the training set size M satisfies $M \geq \frac{128 \ln(\frac{2}{\delta}) (\ln(\max_{\mathcal{J} \subseteq [N^P]} \mathcal{D}_{\mathcal{J}}))^4}{\|\mathcal{D}_{\text{pop}}^P\|^2 \gamma^4}$, where $\mathcal{D}_{\mathcal{J}} :=$
 747 $\min\{\prod_{n \in \mathcal{J}} D_n, \prod_{n \in \mathcal{J}^c} D_n\}$ for $\mathcal{J} \subseteq [N^P]$, then with probability at least $1 - \delta$:

$$|QE(\mathcal{D}_{\text{emp}}^P; \mathcal{J}) - QE(\mathcal{D}_{\text{pop}}^P; \mathcal{J})| \leq \gamma \text{ for all } \mathcal{J} \subseteq [N^P].$$

748 *Proof.* The claim is established by following steps identical to those in the proof of Proposition 1. \square

749 **Corollary 4.** Consider the setting and notation of Corollary 3, with $\epsilon \neq 0$. For $\delta \in (0, 1)$, suppose
 750 that the training set size M satisfies $M \geq \frac{128 \ln(\frac{2}{\delta}) ((16N^P - 24) \ln(\max_{\mathcal{J} \subseteq [N^P]} \mathcal{D}_{\mathcal{J}}))^4}{\|\mathcal{D}_{\text{pop}}^P\|^2 (\ln(R) \cdot \epsilon^2)^4}$. Then, with
 751 probability at least $1 - \delta$ the following hold. First, if there exists a canonical partition $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N^P$
 752 (Definition 6) under which $QE(\mathcal{D}_{\text{emp}}^P; \mu(\mathcal{K})) > (1 + \frac{\epsilon^2}{16N^P - 24}) \cdot \ln(R) + 2\epsilon \cdot \ln(D_{\mu(\mathcal{K})}) + 2\sqrt{2}\epsilon$,
 753 then:

$$\text{SubOpt}^P > \epsilon.$$

754 Second, if for all $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N^P$ it holds that $QE(\mathcal{D}_{\text{emp}}^P; \mu(\mathcal{K})) \leq \frac{\epsilon^2}{16N^P - 24} \cdot \ln(R)$, then:

$$\text{SubOpt}^P \leq \epsilon.$$

755 Moreover, the conditions above on the entanglements of $\mathcal{D}_{\text{emp}}^P$ can be evaluated efficiently (in
 756 $\mathcal{O}(DN^{2P}M^2 + N^P M^3)$ time $\mathcal{O}(DN^P M + M^2)$ and memory, where $D := \max_{n \in [N^P]} D_n$).

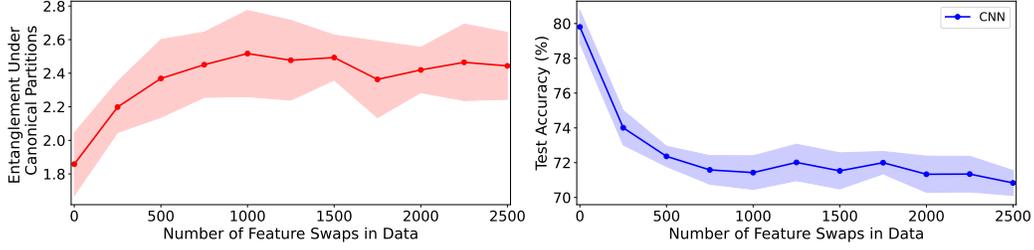


Figure 8: The prediction accuracies of convolutional neural networks are inversely correlated with the entanglements of image data under canonical partitions of features, in compliance with our theory (Appendices I.1 and I.2). This figure is identical to Figure 3, except that the measurements were carried over a binary classification version of the CIFAR10 image dataset, as opposed to a one-dimensional (sequential) audio dataset. For further details see caption of Figure 3 as well as Appendix J.2.

757 *Proof.* Implied by Corollary 3, Proposition 4 with $\gamma = \frac{\epsilon^2}{16N^P - 24} \cdot \ln(R)$ and Algorithm 1 in Ap-
 758 pendix E. \square

759 I.2.3 Empirical Demonstration

760 Figure 8 extends the experiments reported by Figure 3 in Section 4.3 from one-dimensional (se-
 761 quential) audio data to two-dimensional image data. Specifically, it demonstrates that the prediction
 762 accuracies of convolutional neural networks over a variant of CIFAR10 [33] is inversely correlated
 763 with the entanglements that the data admits under canonical partitions of features (Definition 6),
 764 *i.e.* with the entanglements of the empirical data tensor under partitions of its axes induced by
 765 canonical partitions.

766 I.3 Enhancing Suitability of Data to Locally Connected Neural Networks

767 We extend the preprocessing algorithm from Section 5, aimed to enhance the suitability of a data
 768 distribution to locally connected neural networks, from one-dimensional to P -dimensional models
 769 and data (Appendices I.3.1 and I.3.2). Empirical evaluations demonstrate that it significantly improves
 770 prediction accuracy of common architectures (Appendix I.3.3).

771 I.3.1 Search for Feature Arrangement With Low Entanglement Under Canonical Partitions

772 Suppose we have $M \in \mathbb{N}$ training instances $\{\{\mathbf{x}^{((n_1, \dots, n_P), m)}\}_{n_1, \dots, n_P \in [N]}, y^{(m)}\}_{m=1}^M$, where
 773 $\mathbf{x}^{((n_1, \dots, n_P), m)} \in \mathbb{R}^D$ and $y^{(m)} \in \{1, -1\}$ for $n_1, \dots, n_P \in [N], m \in [M]$, with $D \in \mathbb{N}$. For
 774 models intaking P -dimensional data, the recipe for enhancing the suitability of a data distribution to
 775 locally connected neural networks from Section 5.1 boils down to finding a permutation $\pi : [N]^P \rightarrow$
 776 $[N]^P$, which when applied to feature coordinates leads the empirical data tensor $\mathcal{D}_{\text{emp}}^P$ (Equation (7))
 777 to admit low entanglement under canonical partitions (Definition 6).⁸

778 A greedy realization, analogous to that outlined in Section 5.1, is as follows. Initially, partition the
 779 features into 2^P equally sized disjoint sets $\{\mathcal{K}_{1, (k_1, \dots, k_P)} \subset [N]^P\}_{k_1, \dots, k_P \in [2]}$ such that the average
 780 of $\mathcal{D}_{\text{emp}}^P$'s entanglements induced by these sets is minimal. That is, find an element of:

$$\begin{aligned} & \operatorname{argmin} && \frac{1}{2^P} \sum_{k_1, \dots, k_P \in [2]} QE(\mathcal{D}_{\text{emp}}^P; \mu(\mathcal{K}'_{k_1, \dots, k_P})), \\ & \{\mathcal{K}'_{k_1, \dots, k_P} \subset [N]^P\}_{k_1, \dots, k_P \in [2]} && \\ & \text{s.t. } \cup_{k_1, \dots, k_P \in [2]} \mathcal{K}'_{k_1, \dots, k_P} = [N]^P, && \\ & \forall k_1, \dots, k_P, k'_1, \dots, k'_P \in [2] \quad |\mathcal{K}'_{k_1, \dots, k_P}| = |\mathcal{K}'_{k'_1, \dots, k'_P}| && \end{aligned}$$

781 where $\mu : [N]^P \rightarrow [N]^P$ is a compatible map from coordinates in $[N]^P$ to axes indices (Definition 5).
 782 The permutation π will map each $\mathcal{K}_{1, (k_1, \dots, k_P)}$ to coordinates $\times_{p=1}^P \{\frac{N}{2} \cdot (k_p - 1) + 1, \dots, \frac{N}{2} \cdot k_p\}$,
 783 for $k_1, \dots, k_P \in [2]$. Then, partition similarly each of $\{\mathcal{K}_{1, (k_1, \dots, k_P)}\}_{k_1, \dots, k_P \in [2]}$ into 2^P equally

⁸Enhancing the suitability of a data distribution with instances of dimension different than P to P -dimensional models is possible by first arbitrarily mapping features to coordinates in $[N]^P$, and then following the scheme for rearranging P -dimensional data.

Algorithm 3 Enhancing Suitability of P -Dimensional Data to Locally Connected Neural Networks

1: **Input:** $\mathcal{X} := \{\{\mathbf{x}^{((n_1, \dots, n_P), m)}\}_{n_1, \dots, n_P \in [N]}\}_{m=1}^M$ — $M \in \mathbb{N}$ data instances comprising N^P features each

2: **Output:** Permutation $\pi : [N]^P \rightarrow [N]^P$ to apply to feature coordinates

3: Let $\mathcal{K}_{0, (1, \dots, 1)} := [N]^P$ and denote $L := \log_2(N)$

4: # We assume for simplicity that N is a power of two, otherwise one may add constant features

5: **for** $l = 0, \dots, L - 1$, $k_1, \dots, k_P = 1, \dots, 2^l$ **do**

6: Using a reduction to a minimum balanced 2^P -cut problem (Proposition 5), find an approximate solution $\{\mathcal{K}_{l+1, (2k_1-i_1, \dots, 2k_P-i_P)} \subset \mathcal{K}_{l, (k_1, \dots, k_P)}\}_{i_1, \dots, i_P \in \{0, 1\}}$ **for**:

$$\begin{aligned} & \min_{\{\mathcal{K}_{k'_1, \dots, k'_P} \subset \mathcal{K}_{l, (k_1, \dots, k_P)}\}_{k'_1, \dots, k'_P \in [2]}} \frac{1}{2^P} \sum_{k'_1, \dots, k'_P \in [2]} SE(\mathcal{X}; \mathcal{K}_{k'_1, \dots, k'_P}) \\ & \text{s.t. } \cup_{k'_1, \dots, k'_P \in [2]} \mathcal{K}_{k'_1, \dots, k'_P} = \mathcal{K}_{l, (k_1, \dots, k_P)}, \\ & \forall k'_1, \dots, k'_P, \hat{k}_1, \dots, \hat{k}_P \in [2] \quad |\mathcal{K}_{k'_1, \dots, k'_P}| = |\mathcal{K}_{\hat{k}_1, \dots, \hat{k}_P}| \end{aligned}$$

7: **end for**

8: # At this point, $\{\mathcal{K}_{L, (k_1, \dots, k_P)}\}_{k_1, \dots, k_P \in [N]}$ each contain a single feature coordinate

9: **return** π that maps $(n_1, \dots, n_P) \in \mathcal{K}_{L, (k_1, \dots, k_P)}$ to (k_1, \dots, k_P) , for every $k_1, \dots, k_P \in [N]$

784 sized disjoint sets. Continuing in the same fashion, until we reach subsets $\{\mathcal{K}_{L, (k_1, \dots, k_P)}\}_{k_1, \dots, k_P \in [N]}$
 785 consisting of a single feature coordinate each, fully specifies the permutation π .

786 As in the case of one-dimensional models and data (Section 5.1), the step lying at the heart of the
 787 above scheme — finding a balanced partition into 2^P sets that minimizes average entanglement — is
 788 computationally prohibitive. In the next supappendix we will see that replacing entanglement with
 789 surrogate entanglement brings forth a practical implementation.

790 I.3.2 Practical Algorithm via Surrogate for Entanglement

791 To efficiently implement the scheme from Appendix I.3.1, we replace entanglement with surrogate
 792 entanglement (Definition 4), which for P -dimensional data is straightforwardly defined as follows.

793 **Definition 8.** Given a set of $M \in \mathbb{N}$ instances $\mathcal{X} := \{\{\mathbf{x}^{((n_1, \dots, n_P), m)} \in \mathbb{R}^D\}_{n_1, \dots, n_P \in [N]}\}_{m=1}^M$,
 794 denote by $p_{(n_1, \dots, n_P), (n'_1, \dots, n'_P)}$ the multivariate Pearson correlation between features $(n_1, \dots, n_P) \in$
 795 $[N]^P$ and $(n'_1, \dots, n'_P) \in [N]^P$. For $\mathcal{K} \subseteq [N]^P$, the *surrogate entanglement* of \mathcal{X} with respect to
 796 the partition $(\mathcal{K}, \mathcal{K}^c)$, denoted $SE(\mathcal{X}; \mathcal{K})$, is the sum of Pearson correlation coefficients between
 797 pairs of features, the first belonging to \mathcal{K} and the second to $\mathcal{K}^c := [N]^P \setminus \mathcal{K}$. That is:

$$SE(\mathcal{X}; \mathcal{K}) = \sum_{(n_1, \dots, n_P) \in \mathcal{K}, (n'_1, \dots, n'_P) \in \mathcal{K}^c} p_{(n_1, \dots, n_P), (n'_1, \dots, n'_P)}.$$

798 Analogously to the case of one-dimensional data (Appendix F), Proposition 5 shows that replacing
 799 the entanglement with surrogate entanglement in the scheme from Appendix I.3.1 converts each
 800 search for a balanced partition minimizing average entanglement into a *minimum balanced 2^P -cut*
 801 *problem*. Although the minimum balanced 2^P -cut problem is NP-hard (see, e.g., [23]), similarly to
 802 the minimum balanced (2-) cut problem, it enjoys a wide array of well-established approximation
 803 implementations, particularly ones designed for large scale [29, 56]. We therefore obtain a practical
 804 algorithm for enhancing the suitability of a data distribution with P -dimensional instances to locally
 805 connected neural networks — see Algorithm 3.

806 **Proposition 5.** For any $\bar{\mathcal{K}} \subseteq [N]^P$ of size divisible by 2^P , the following optimization problem can
 807 be framed as a minimum balanced 2^P -cut problem over a complete graph with $|\bar{\mathcal{K}}|$ vertices:

$$\begin{aligned} & \min_{\{\mathcal{K}_{k_1, \dots, k_P} \subset \bar{\mathcal{K}}\}_{k_1, \dots, k_P \in [2]}} \frac{1}{2^P} \sum_{k_1, \dots, k_P \in [2]} SE(\mathcal{X}; \mathcal{K}_{k_1, \dots, k_P}). \quad (9) \\ & \text{s.t. } \cup_{k_1, \dots, k_P \in [2]} \mathcal{K}_{k_1, \dots, k_P} = \bar{\mathcal{K}}, \\ & \forall k_1, \dots, k_P, k'_1, \dots, k'_P \in [2] \quad |\mathcal{K}_{k_1, \dots, k_P}| = |\mathcal{K}_{k'_1, \dots, k'_P}| \end{aligned}$$

Table 3: Arranging features of randomly permuted image data via Algorithm 3 significantly improves the prediction accuracies of convolutional neural networks. Reported are the results of experiments analogous to those of Table 1, carried out over the CIFAR10 image dataset (as opposed to an audio dataset) using Algorithm 3. For further details see caption of Table 1 as well as Appendix J.2.

	Randomly Permuted	Algorithm 3	IGTD
CNN	35.1 \pm 0.5	38.2 \pm 0.4	36.2 \pm 0.7

Table 4: Our feature rearrangement method (Algorithm 2 in Appendix F) can be efficiently applied to data with a large number of features. Reported are the results of an experiment identical to that of Table 1, but over a version of the Speech Commands [62] dataset in which every instance has 50,000 features. Instances of the minimum balanced cut problem solved as part of Algorithm 2 entail graphs with up to $25 \cdot 10^8$ edges. They are solved using the well-known edge sparsification algorithm of [56] that preserves weights of cuts, allowing for configurable compute and memory consumption (the more resources are consumed, the more accurate the solution will be). As can be seen, Algorithm 2 leads to significant improvements in prediction accuracies. See Appendix J.2 for further implementation details.

	Randomly Permuted	Our Method
CNN	15.0 \pm 1.6	65.6 \pm 1.1
S4	18.2 \pm 0.5	82.2 \pm 0.4

808 Specifically, there exists a complete undirected weighted graph with vertices $\bar{\mathcal{K}}$ and edge
809 weights $w : \bar{\mathcal{K}} \times \bar{\mathcal{K}} \rightarrow \mathbb{R}$ such that, for any partition of $\bar{\mathcal{K}}$ into equally sized dis-
810 joint sets $\{\mathcal{K}_{k_1, \dots, k_P}\}_{k_1, \dots, k_P \in [2]}$, the weight of the 2^P -cut in the graph induced by them —
811 $\frac{1}{2} \sum_{k_1, \dots, k_P \in [2]} \sum_{(n_1, \dots, n_P) \in \mathcal{K}_{k_1, \dots, k_P}, (n'_1, \dots, n'_P) \in \bar{\mathcal{K}} \setminus \mathcal{K}_{k_1, \dots, k_P}} w(\{(n_1, \dots, n_P), (n'_1, \dots, n'_P)\})$ —
812 is equal, up to multiplicative and additive constants, to the term minimized in Equation (4).

813 *Proof.* Consider the complete undirected graph whose vertices are $\bar{\mathcal{K}}$ and where the weight of an
814 edge $\{(n_1, \dots, n_P), (n'_1, \dots, n'_P)\} \in \bar{\mathcal{K}} \times \bar{\mathcal{K}}$ is:

$$w(\{(n_1, \dots, n_P), (n'_1, \dots, n'_P)\}) = p_{(n_1, \dots, n_P), (n'_1, \dots, n'_P)}$$

815 (recall that $p_{(n_1, \dots, n_P), (n'_1, \dots, n'_P)}$ stands for the multivariate Pearson correlation between features
816 (n_1, \dots, n_P) and (n'_1, \dots, n'_P) in \mathcal{X}). For any partition of $\bar{\mathcal{K}}$ into equally sized disjoint sets
817 $\{\mathcal{K}_{k_1, \dots, k_P}\}_{k_1, \dots, k_P \in [2]}$, it holds that:

$$\begin{aligned} & \frac{1}{2} \sum_{k_1, \dots, k_P \in [2]} \sum_{(n_1, \dots, n_P) \in \mathcal{K}_{k_1, \dots, k_P}, (n'_1, \dots, n'_P) \in \bar{\mathcal{K}} \setminus \mathcal{K}_{k_1, \dots, k_P}} w(\{(n_1, \dots, n_P), (n'_1, \dots, n'_P)\}) \\ &= \frac{1}{2} \sum_{k_1, \dots, k_P \in [2]} SE(\mathcal{X}; \mathcal{K}_{k_1, \dots, k_P}) - \frac{1}{2} SE(\mathcal{X}; \bar{\mathcal{K}}), \end{aligned}$$

818 where $\frac{1}{2} SE(\mathcal{X}; \bar{\mathcal{K}})$ does not depend on $\{\mathcal{K}_{k_1, \dots, k_P}\}_{k_1, \dots, k_P \in [2]}$. This concludes the proof. \square

819 I.3.3 Experiments

820 We supplement the experiments from Section 5.2 for one-dimensional models and data, by empiri-
821 cally evaluating Algorithm 3 using two-dimensional convolutional neural networks over randomly
822 permuted image data. Specifically, Table 3 presents experiments analogous to those of Table 1
823 from Section 5.2.1 over the CIFAR10 [33] image dataset. For brevity, we defer some implementation
824 details to Appendix J.2.

825 J Further Experiments and Implementation Details

826 J.1 Further Experiments

827 The number of features in the audio dataset used in the experiment of Section 5.2.1 is 2048. We
828 demonstrate the scalability of our feature rearrangement method (Algorithm 2 in Appendix F) by
829 including in Table 4 an experiment over audio data with 50,000 features. In this experiment, instances

Table 5: Arranging features of tabular datasets via our method (detailed in Appendix F) significantly improves the prediction accuracies of locally connected neural networks. Reported are results of experiments identical to those of Table 2, but over the “dna” tabular classification dataset [59]. For further details, see caption of Table 2 and Appendix J.2.

	Dataset: dna		
	Baseline	Our Method	IGTD
CNN	82.5 ± 1.7	91.2 ± 1.1	87.4 ± 1.1
S4	86.4 ± 1.7	89.1 ± 3.7	89.9 ± 1.1
Local-Attention	79.2 ± 4.0	85.7 ± 4.5	82.7 ± 3.2

830 of the minimum balanced cut problem encountered in Algorithm 2 entail graphs with up to $25 \cdot 10^8$
831 edges. They are solved using the well-known edge sparsification algorithm of [56] that preserves
832 weights of cuts, allowing for configurable compute and memory consumption (the more resources are
833 consumed, the more accurate the solution will be). In contrast, this approach is not directly applicable
834 for improving the efficiency of IGTD.

835 Table 5 supplements Table 2 from Section 5.2.2 by reporting results of experiments with an additional
836 tabular benchmark — “dna”.

837 J.2 Further Implementation Details

838 We provide implementation details omitted from our experimental reports (Section 4.3, Section 5, Ap-
839 pendix I.2.3, Appendix I.3.3 and Appendix J.1). We provide implementation details omitted from our
840 experimental reports (Section 4.3, Section 5 and Appendix I.2.3). Source code for reproducing our
841 results and figures, based on the PyTorch [44] framework, is attached as supplementary material and
842 will be made publicly available. All experiments were run on a single Nvidia RTX A6000 GPU.

843 J.3 Empirical Demonstration of Theoretical Analysis (Figures 3 and 8)

844 J.3.1 Figure 3

845 **Dataset:** The SpeechCommands dataset [62] contains raw audio segments of length up to 16000,
846 split into 84843 train and 11005 test segments. We zero-padded all audio segments to have a length of
847 16000 and resampled them using sinc interpolation (default PyTorch implementation). We allocated
848 11005 audio segments from the train set for validation, *i.e.* the dataset was split into 73838 train,
849 11005 validation and 11005 test audio segments, and created a binary one-vs-all classification version
850 of SpeechCommands by taking all audio segments labeled by the class “33” (corresponding to the
851 word “yes”), and sampling an equal amount of segments from the remaining classes (this process was
852 done separately for the train, validation and test sets). The resulting balanced classification dataset
853 had 5610 train, 846 validation and 838 test segments. Lastly, we resampled all audio segments in the
854 dataset from 16000 to 4096 features using sinc interpolation.

855 **Random feature swaps:** Starting with the original order of features, we created increasingly “shuf-
856 fled” versions of the dataset by randomly swapping the position of features. For each number of
857 random position swaps $k \in \{0, 250, \dots, 2500\}$, we created ten datasets, whose features were subject
858 to k random position swaps between features, using different random seeds.

859 **Quantum entanglement measurement:** Each reported value in the plot is the average of entangle-
860 ments with respect to canonical partitions (Definition 2) corresponding to levels $l = 1, 2, 3, 4, 5, 6$.
861 We used Algorithm 1 described in Appendix E on two mini-batches of size 500, randomly sampled
862 from the train set. As customary (*cf.* [58]), every input feature x was embedded using the following
863 sine-cosine scheme:

$$\phi(x) := (\sin(\pi\theta x), \cos(\pi\theta x)) \in \mathbb{R}^2,$$

864 with $\theta = 0.085$.

865 Neural network architectures:

- 866 • **CNN:** An adaptation of the M5 architecture from [16], which is designed for audio data. Our
867 implementation is based on [https://github.com/danielajisafe/Audio_WaveForm_](https://github.com/danielajisafe/Audio_WaveForm_Paper_Implementation)
868 [Paper_Implementation](https://github.com/danielajisafe/Audio_WaveForm_Paper_Implementation).
- 869 • **S4:** Official implementation of [26] with a hidden dimension of 128 and 4 layers.

870 • **Local-Attention:** Adaptation of the local-attention model from [46], as implemented in
 871 <https://github.com/lucidrains/local-attention>. We use a multi-layer percep-
 872 tron (MLP) for mapping continuous (raw) inputs to embeddings, which are fed as input
 873 the the local-attention model. For classification, we collapse the spatial dimension of the
 874 network’s output using mean-pooling and pass the result into an MLP classification head.
 875 The network had attention dimension 128 (with 2 heads of dimension 64), depth 8, hidden
 876 dimension of MLP blocks 341 (computed automatically by the library based on the attention
 877 dimension) and local-attention window size 10.

878 **Training and evaluation:** The binary cross-entropy loss was minimized via the Adam optimizer [32]
 879 with default β_1, β_2 coefficients. Batch sizes were chosen to be the largest powers of two that fit in
 880 the GPU memory. The duration of optimization was 150, 40 and 50 epochs for the CNN, S4 and
 881 Local-Attention models, respectively (number of epochs was taken large enough such that the train
 882 loss plateaued). After the last training epoch, the model which performed best on the validation set
 883 was chosen, and its average test accuracy is reported. Additional optimization hyperparameters are
 884 provided in Table 6. We note that for S4 [26], in accordance with its official implementation, we use
 885 a cosine annealing learning rate scheduler.

Table 6: Optimization hyperparameters for the experiments of Figure 3.

Model	Optimizer	Learning Rate	Weight Decay	Batch Size
CNN	Adam	0.01	0.0001	128
S4	AdamW	0.001	0.01	64
Local-Attention	Adam	0.0005	0	32

886 **J.3.2 Figure 8**

887 **Dataset:** We created one-vs-all binary classification datasets based on CIFAR10 [33] as follows. All
 888 images were converted to grayscale using the PyTorch default implementation. Then, we allocated
 889 7469 images from the train set for validation, *i.e.* the dataset was split into 42531 train, 7469 validation
 890 and 1000 test images. We took all images labeled by the class “0” (corresponding to images of
 891 airplanes), and uniformly sampled an equal amount of images from the remaining classes (this process
 892 was done separately for the train, validation and test sets). The resulting balanced classification
 893 dataset had 8506 train, 1493 validation and 2001 test images.

894 **Random feature swaps:** We created increasingly “shuffled” versions of the dataset according to the
 895 protocol described in Appendix J.3.1.

896 **Quantum entanglement measurement:** Each reported value in the plot is the average entanglements
 897 with respect to canonical partitions (Definition 6) corresponding to levels $l = 1, 2$.

898 **Neural network architectures:**

899 • **CNN:** Same architecture used for the experiments of Table 2, but with one-dimensional
 900 convolutional layers replaced with two-dimensional convolutional layers. See Table 7 for
 901 the exact architectural hyperparameters.

Table 7: Architectural hyperparameters for the convolutional neural network used in the experiments of Figure 8 and Table 3.

Hyperparameter	Value
Stride	(3, 3)
Kernel size	(3, 3)
Pooling window size	(3, 3)
Number of blocks	8
Hidden dimension	32

902 **Training and evaluation:** The binary cross-entropy loss was minimized for 150 epochs via the
 903 Adam optimizer [32] with default β_1, β_2 coefficients (number of epochs was taken large enough such
 904 that the train loss plateaued). Batch size was chosen to be the largest power of two that fit in the

905 GPU memory. After the last training epoch, the model which performed best on the validation set
 906 was chosen, and its average test accuracy is reported. Additional optimization hyperparameters are
 907 provided in Table 8.

Table 8: Optimization hyperparameters for the experiments of Figure 8.

Model	Optimizer	Learning Rate	Weight Decay	Batch Size
CNN	Adam	0.001	0.0001	128

908 **J.4 Enhancing Suitability of Data to Locally Connected Neural Networks (Tables 1, 2, 3,
 909 and 4)**

910 **J.4.1 Features Rearrangement Algorithms**

911 Algorithm 2, 3 and IGTD [65] were applied to the training set. Subsequently, the learned feature
 912 rearrangement was used for both the validation and test data. In instances where three-way cross-
 913 validation was used, distinct rearrangements were learned for each separately.

914 **Algorithm 2 and Algorithm 3:** Approximate solutions to the minimum balanced cut problems were
 915 obtained using the METIS graph partitioning algorithm [29], as implemented in <https://github.com/networkx/networkx-metis>.
 916

917 **IGTD [65]:** The original IGTD implementation supports rearranging data only into two-dimensional
 918 images. We adapted its implementation to support one-dimensional (sequential) data for the experi-
 919 ments of Tables 1 and 2.

920 **J.4.2 Randomly Permuted Audio Datasets (Table 1)**

921 **Dataset:** To facilitate efficient experimentation, we downsampled all audio segments in SpeechCom-
 922 mands to have 2048 features. Furthermore, for the train, validation and test sets separately, we used
 923 20% of the audio segments available for each class.

924 **Neural network architectures:**

- 925 • **CNN:** Same architecture used for the experiments of Figure 3 (see Appendix J.3.1).
- 926 • **S4:** Same architecture used for the experiments of Figure 3 (see Appendix J.3.1).
- 927 • **Local-Attention:** Same architecture used in the experiments of Figure 3, but with a depth 4
 928 network (we reduced the depth to allow for more efficient experimentation).

929 **Training and evaluation:** The cross-entropy loss was minimized via the Adam optimizer [32] with
 930 default β_1, β_2 coefficients. Batch sizes were chosen to be the largest powers of two that fit in the
 931 GPU memory. The duration of optimization was 200, 200 and 450 epochs for the CNN, S4 and
 932 Local-Attention models, respectively (number of epochs was taken large enough such that the train
 933 loss plateaued). After the last training epoch, the model which performed best on the validation set
 934 was chosen, and its average test accuracy is reported. Additional optimization hyperparameters are
 935 provided in Table 9. We note that for S4 [26], in accordance with its official implementation, we use
 936 a cosine annealing learning rate scheduler.

Table 9: Optimization hyperparameters for the experiments of Table 1.

Model	Optimizer	Learning Rate	Weight Decay	Batch Size
CNN	Adam	0.001	0.0001	128
S4	AdamW	0.001	0.01	64
Local-Attention	Adam	0.0001	0	32

937 **J.4.3 Tabular Datasets (Tables 2 and 5)**

938 **Datasets:** The datasets "dna", "semeion" and "isolet" are all from the OpenML repository [59]. For
 939 each dataset we split the samples into three folds, which were used for evaluation according to a
 940 standard three-way cross-validation protocol. That is, for each of the three folds, we used one third

941 of the data as a test set and the remaining for train and validation. One third of the samples in the
 942 remaining folds (not used for testing) were allocated for the validation set.

943 **Neural network architectures:**

- 944 • **CNN:** We used a ResNet adapted for tabular data. It consisted of residual blocks of the
 945 following form:

$$\text{Block}(\mathbf{x}) = \text{dropout}(\mathbf{x} + \text{BN}(\text{maxpool}(\text{ReLU}(\text{conv}(\mathbf{x}))))).$$

946 After applying a predetermined amount of residual blocks, a global average pooling and
 947 fully connected layers were used to output the prediction. The architectural hyperparameters
 948 are specified in Table 10.

- 949 • **S4:** Same architecture used for the experiments of Figure 3 (see Appendix J.3.1), but with a
 950 hidden dimension of 64.
- 951 • **Local-Attention:** Same architecture used for the experiments of Figure 3 (see Ap-
 952 pendix J.3.1), but with 4 attention heads of dimension 32 and a local-attention window size
 953 of 25.

Table 10: Architectural hyperparameters for the convolutional neural network used in the experiments of Tables 2 and 4.

Hyperparameter	Value
Stride	3
Kernel size	3
Pooling window size	3
Number of blocks	8
Hidden dimension	32

954 **Training and evaluation:** The cross-entropy loss was minimized for 300 epochs via the Adam
 955 optimizer [32] with default β_1, β_2 coefficients (number of epochs was taken large enough such that
 956 the train loss plateaued). Batch sizes were chosen to be the largest powers of two that fit in the GPU
 957 memory. After the last training epoch, the model which performed best according to the validation
 958 sets was chosen, and test accuracy was measured on the test set. The reported accuracy is the average
 959 over the three folds. Additional optimization hyperparameters are specified in Table 11. We note that
 960 for S4 [26], in accordance with its official implementation, we use a cosine annealing learning rate
 961 scheduler.

Table 11: Optimization hyperparameters for the experiments of Table 2.

Model	Optimizer	Learning Rate	Weight Decay	Batch Size
CNN	Adam	0.001	0.0001	64
S4	AdamW	0.001	0.01	64
Local-Attention	Adam	0.00005	0	64

962 **J.4.4 Randomly Permuted Image Datasets (Table 3)**

963 **Dataset:** The data acquisition process followed the protocol described in Appendix J.3.2, except that
 964 the data was not converted into a binary one-vs-all classification dataset.

965 **Neural network architectures:**

- 966 • **CNN:** Same architecture used for the experiments of Figure 8.

967 **Training and evaluation:** The cross-entropy loss was minimized for 500 epochs via the Adam
 968 optimizer [32] with default β_1, β_2 coefficients (number of epochs was taken large enough such
 969 that the train loss plateaued). Batch size was chosen to be the largest power of two that fit in the
 970 GPU memory. After the last training epoch, the model which performed best on the validation set
 971 was chosen, and its average test accuracy is reported. Additional optimization hyperparameters are
 972 provided in Table 12.

Table 12: Optimization hyperparameters for the experiments of Table 3.

Model	Optimizer	Learning Rate	Weight Decay	Batch Size
CNN	Adam	0.001 (multiplied by 0.1 after 300 epochs)	0.0001	128

973 J.4.5 Randomly Permuted Audio Datasets With a Large Number of Features (Table 4)

974 **Dataset:** The data acquisition process followed the protocol described in Appendix J.3.1, except that
 975 the data was not transformed into a binary one-vs-all classification dataset and the audio segments
 976 were upsampled from 16,000 to 50,000.

977 **Edge Sparsification:** To facilitate running the METIS graph partitioning algorithm over the minimum
 978 balanced cut problems encountered as part of Algorithm 2, we first removed edges from the graph us-
 979 ing the spectral sparsification algorithm of [56]. Specifically, we used the official Julia implementation
 980 (<https://github.com/danspielman/Laplacians.jl>) with hyperparameter $\epsilon = 0.15$.

981 **Neural network architectures:**

- 982 • **CNN:** Same architecture used for the experiments of Table 2 (see Appendix J.4.3).
- 983 • **S4:** Same architecture used for the experiments of Table 1 (see Appendix J.3.1), but
 984 with a hidden dimension of 32 (we reduced the hidden dimension due to GPU memory
 985 considerations).

986 **Training and evaluation:** The cross-entropy loss was minimized for 200 epochs via the Adam
 987 optimizer [32] with default β_1, β_2 coefficients (number of epochs was taken large enough such that
 988 the train loss plateaued). Batch sizes were chosen to be the largest powers of two that fit in the
 989 GPU memory. After the last training epoch, the model which performed best on the validation set
 990 was chosen, and its average test accuracy is reported. Additional optimization hyperparameters are
 991 provided in Table 13. We note that for S4 [26], in accordance with its official implementation, we use
 992 a cosine annealing learning rate scheduler.

Table 13: Optimization hyperparameters for the experiments of Table 4.

Model	Optimizer	Learning Rate	Weight Decay	Batch Size
CNN	Adam	0.001	0.0001	64
S4	AdamW	0.001	0.01	64

993 K Deferred Proofs

994 K.1 Useful Lemmas

995 Below we collect a few useful results, which we will use in our proofs.

996 **Lemma 1.** *We denote the vector of singular values of a matrix \mathbf{X} (arranged in decreasing order) by*
 997 *$S(\mathbf{X})$. For any $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{D_1 \times D_2}$ it holds that:*

$$\|S(\mathbf{A}) - S(\mathbf{B})\| \leq \|S(\mathbf{A} - \mathbf{B})\|$$

998 *Proof.* See Theorem III.4.4 of [3]. □

999 **Lemma 2.** *Let $\mathcal{P} = \{p_1, \dots, p_N\}, \mathcal{Q} = \{q_1, \dots, q_N\}$ be two probability distributions supported on*
 1000 *$[N]$, and denote by $TV(\mathcal{P}, \mathcal{Q}) := \frac{1}{2} \sum_{n=1}^N |p_n - q_n|$ their total variation distance. If for $\epsilon \in (0, 1/2)$*
 1001 *it holds that $TV(\mathcal{P}, \mathcal{Q}) \leq \epsilon$, then:*

$$|H(\mathcal{P}) - H(\mathcal{Q})| \leq H_b(\epsilon) + \epsilon \cdot \ln(N),$$

1002 *where $H(\mathcal{P}) := -\sum_{n=1}^N p_n \ln(p_n)$ is the entropy of \mathcal{P} , and $H_b(c) := -c \cdot \ln(c) - (1-c) \cdot \ln(1-c)$*
 1003 *is the binary entropy of a Bernoulli distribution parameterized by $c \in [0, 1]$.*

1004 *Proof.* See, e.g., Theorem 11 of [27]. □

1005 **Lemma 3** (Hoeffding inequality in Hilbert space). *Let X_1, \dots, X_N be an i.i.d sequence of random*
 1006 *variables whose range is some separable Hilbert space \mathcal{H} . Suppose that $\mathbb{E}[X_n] = 0$ and $\|X_n\| \leq c$*
 1007 *for all $n \in [N]$. Then, for all $t \geq 0$:*

$$Pr \left(\left\| \frac{1}{n} \sum_{n=1}^N X_n \right\| \geq t \right) \leq 2 \exp \left(-\frac{Nt^2}{2c^2} \right),$$

1008 where $\|\cdot\|$ refers to the Hilbert space norm.

1009 *Proof.* See Section 2.4 of [51]. □

1010 **Lemma 4** (adapted from [35]). *Let $G = (V, E)$ be the perfect binary tree graph, with vertices V*
 1011 *and edges E , underlying the locally connected tensor network that generates $\mathcal{W}_{\text{TN}} \in \mathbb{R}^{D_1 \times \dots \times D_N}$*
 1012 *(defined in Section 3.1). For $\mathcal{K} \subseteq [N]$, let $V_{\mathcal{K}} \subseteq V$ and $V_{\mathcal{K}^c} \subseteq V$ be the leaves in G corresponding to*
 1013 *axes indices \mathcal{K} and \mathcal{K}^c of \mathcal{W}_{TN} , respectively. Lastly, given a cut (A, B) of V , i.e. $A \subseteq V$ and $B \subseteq V$*
 1014 *are disjoint and $A \cup B = V$, denote by $C(A, B) := \{\{u, v\} \in E : u \in A, v \in B\}$ the edge-cut set.*
 1015 *Then:*

$$\text{rank}(\llbracket \mathcal{W}_{\text{TN}}; \mathcal{K} \rrbracket) \leq \min_{\text{cut } (A, B) \text{ of } V \text{ s.t. } V_{\mathcal{K}} \subseteq A, V_{\mathcal{K}^c} \subseteq B} R^{|C(A, B)|}.$$

1016 *In particular, if $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$, then:*

$$\text{rank}(\llbracket \mathcal{W}_{\text{TN}}; \mathcal{K} \rrbracket) \leq R.$$

1017 *Proof.* See Claim 1 in [35] for the upper bound on the rank of $\llbracket \mathcal{W}_{\text{TN}}; \mathcal{K} \rrbracket$ for any $\mathcal{K} \subseteq [N]$. If
 1018 $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$, since there exists a cut (A, B) of V such that $V_{\mathcal{K}} \subseteq A$ and $V_{\mathcal{K}^c} \subseteq B$ whose edge-cut
 1019 set is of a singleton, we get that $\text{rank}(\llbracket \mathcal{W}_{\text{TN}}; \mathcal{K} \rrbracket) \leq R$. □

1020 **Lemma 5** (adapted from [35]). *Let $G = (V, E)$ be the perfect 2^P -ary tree graph, with vertices V*
 1021 *and edges E , underlying the locally connected tensor network that generates $\mathcal{W}_{\text{TN}}^P \in \mathbb{R}^{D_1 \times \dots \times D_{N^P}}$*
 1022 *(defined in Appendix I.1.1). Furthermore, let $\mu : [N]^P \rightarrow [N^P]$ be a compatible map from P -*
 1023 *dimensional coordinates to axes indices of $\mathcal{W}_{\text{TN}}^P$ (Definition 5). For $\mathcal{J} \subseteq [N^P]$, let $V_{\mathcal{J}} \subseteq V$*
 1024 *and $V_{\mathcal{J}^c} \subseteq V$ be the leaves in G corresponding to axes indices \mathcal{J} and \mathcal{J}^c of $\mathcal{W}_{\text{TN}}^P$, respectively.*
 1025 *Lastly, given a cut (A, B) of V , i.e. $A \subseteq V$ and $B \subseteq V$ are disjoint and $A \cup B = V$, denote by*
 1026 *$C(A, B) := \{\{u, v\} \in E : u \in A, v \in B\}$ the edge-cut set. Then:*

$$\text{rank}(\llbracket \mathcal{W}_{\text{TN}}^P; \mathcal{J} \rrbracket) \leq \min_{\text{cut } (A, B) \text{ of } V \text{ s.t. } V_{\mathcal{J}} \subseteq A, V_{\mathcal{J}^c} \subseteq B} R^{|C(A, B)|}.$$

1027 *In particular, if $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N^P$, then:*

$$\text{rank}(\llbracket \mathcal{W}_{\text{TN}}; \mu(\mathcal{K}) \rrbracket) \leq R.$$

1028 *Proof.* See Claim 1 in [35] for the upper bound on the rank of $\llbracket \mathcal{W}_{\text{TN}}^P; \mathcal{J} \rrbracket$ for any $\mathcal{J} \subseteq [N^P]$. If
 1029 $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N^P$, since there exists a cut (A, B) of V such that $V_{\mu(\mathcal{K})} \subseteq A$ and $V_{\mu(\mathcal{K}^c)} \subseteq B$ whose
 1030 edge-cut set is of a singleton, we get that $\text{rank}(\llbracket \mathcal{W}_{\text{TN}}^P; \mu(\mathcal{K}) \rrbracket) \leq R$. □

1031 **Lemma 6** (adapted from Theorem 3.18 in [25]). *Let $\mathcal{A} \in \mathbb{R}^{D_1 \times \dots \times D_N}$ and $\epsilon > 0$. For each*
 1032 *canonical partition $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$, let $\sigma_{\mathcal{K}, 1} \geq \dots \geq \sigma_{\mathcal{K}, D_{\mathcal{K}}}$ be the singular values of $\llbracket \mathcal{A}; \mathcal{K} \rrbracket$, where*
 1033 *$D_{\mathcal{K}} := \min\{\prod_{n \in \mathcal{K}} D_n, \prod_{n \in \mathcal{K}^c} D_n\}$. Let $(n_{\mathcal{K}})_{\mathcal{K} \in \mathcal{C}_N} \in \mathbb{N}^{\mathcal{C}_N}$ be an assignment of an integer to each*
 1034 *$\mathcal{K} \in \mathcal{C}_N$. For any such assignment, consider the set of tensors with Hierarchical Tucker (HT) rank at*
 1035 *most $(n_{\mathcal{K}})_{\mathcal{K} \in \mathcal{C}_N}$ as follows:*

$$HT((n_{\mathcal{K}})_{\mathcal{K} \in \mathcal{C}_N}) := \{ \mathcal{V} \in \mathbb{R}^{D_1 \times \dots \times D_N} : \forall \mathcal{K} \in \mathcal{C}_N, \text{rank}(\llbracket \mathcal{V}; \mathcal{K} \rrbracket) \leq n_{\mathcal{K}} \}.$$

1036 *Suppose that for all $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$ it holds that $\sqrt{\sum_{d=n_{\mathcal{K}}+1}^{D_{\mathcal{K}}} \sigma_{\mathcal{K}, d}^2} \leq \frac{\epsilon}{\sqrt{2N-3}}$. Then, there exists*
 1037 *$\mathcal{W} \in HT((n_{\mathcal{K}})_{\mathcal{K} \in \mathcal{C}_N})$ satisfying:*

$$\|\mathcal{W} - \mathcal{A}\| \leq \epsilon.$$

1038 *In particular, if for all $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$ it holds that $\sqrt{\sum_{d=R+1}^{D_{\mathcal{K}}} \sigma_{\mathcal{K}, d}^2} \leq \frac{\epsilon}{\sqrt{2N-3}}$, then there exists*
 1039 *$\mathcal{W}_{\text{TN}} \in \mathbb{R}^{D_1 \times \dots \times D_N}$ generated by the locally connected tensor network (defined in Section 3.1)*
 1040 *satisfying:*

$$\|\mathcal{W}_{\text{TN}} - \mathcal{A}\| \leq \epsilon.$$

1041 **Lemma 7** (adapted from Theorem 3.18 in [25]). Let $(n_{\mathcal{K}})_{\mathcal{K} \in \mathcal{C}_N} \in \mathbb{N}^{\mathcal{C}_N}$ be an assignment of an
 1042 integer to each $\mathcal{K} \in \mathcal{C}_N$. For any such assignment, consider the set of tensors with Hierarchical
 1043 Tucker (HT) rank at most $(n_{\mathcal{K}})_{\mathcal{K} \in \mathcal{C}_N}$ as follows:

$$HT((n_{\mathcal{K}})_{\mathcal{K} \in \mathcal{C}_N}) := \{\mathcal{V} \in \mathbb{R}^{D_1 \times \dots \times D_N} : \forall \mathcal{K} \in \mathcal{C}_N, \text{rank}(\llbracket \mathcal{V}; \mathcal{K} \rrbracket) \leq n_{\mathcal{K}}\}.$$

1044 Consider a locally connected tensor network with varying widths $(n_{\mathcal{K}})_{\mathcal{K} \in \mathcal{C}_N}$, i.e. a tensor network
 1045 conforming to a perfect binary tree graph in which the lengths of inner axes are as follows (in contrast
 1046 to all being equal to R as in the locally connected tensor network defined in Section 3.1). An axis
 1047 corresponding to an edge that connects a node with descendant leaves indexed by \mathcal{K} to its parent
 1048 is assigned the length $n_{\mathcal{K}}$. Then, every $\mathcal{A} \in HT((n_{\mathcal{K}})_{\mathcal{K} \in \mathcal{C}_N})$ can be represented by said locally
 1049 connected tensor network with varying widths, meaning there exists an assignment to the tensors of
 1050 the tensor network such that it generates \mathcal{A} . In particular, if $n_{\mathcal{K}} = R$ for all $\mathcal{K} \in \mathcal{C}_N$, then \mathcal{A} can be
 1051 generated by the locally connected tensor network with all inner axes being of length R .

1052 **Lemma 8.** Let $\mathcal{V}, \mathcal{W} \in \mathbb{R}^{D_1 \times \dots \times D_N}$ be tensors such that $\|\mathcal{V}\| = \|\mathcal{W}\| = 1$ and $\|\mathcal{V} - \mathcal{W}\| < \frac{1}{2}$.
 1053 Then, for any $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$ it holds that:

$$|QE(\mathcal{V}; \mathcal{K}) - QE(\mathcal{W}; \mathcal{K})| \leq H_b(\|\mathcal{V} - \mathcal{W}\|) + \|\mathcal{V} - \mathcal{W}\| \cdot \ln(D_{\mathcal{K}}),$$

1054 where $H_b(c) := -(c \cdot \ln(c) + (1 - c) \cdot \ln(1 - c))$ is the binary entropy of a Bernoulli distribution
 1055 parameterized by $c \in [0, 1]$, and $D_{\mathcal{K}} := \min\{\prod_{n \in \mathcal{K}} D_n, \prod_{n \in \mathcal{K}^c} D_n\}$.

1056 *Proof.* For any matrix $\mathbf{M} \in \mathbb{R}^{D_1 \times D_2}$ with $D := \min\{D_1, D_2\}$, let $S(\mathbf{M}) = (\sigma_{\mathbf{M},1}, \dots, \sigma_{\mathbf{M},D})$ be
 1057 the vector consisting of its singular values. First note that

$$\|\mathcal{V} - \mathcal{W}\| = \|S(\llbracket \mathcal{V} - \mathcal{W}; \mathcal{K} \rrbracket)\| \leq \|\mathcal{V} - \mathcal{W}\|.$$

1058 So by Lemma 1 we have

$$\|S(\llbracket \mathcal{V}; \mathcal{K} \rrbracket) - S(\llbracket \mathcal{W}; \mathcal{K} \rrbracket)\| \leq \|\mathcal{V} - \mathcal{W}\|.$$

1059 Let $v_1 \geq \dots \geq v_{D_{\mathcal{K}}}$ and $w_1 \geq \dots \geq w_{D_{\mathcal{K}}}$ be the singular values of $\llbracket \mathcal{V}; \mathcal{K} \rrbracket$ and $\llbracket \mathcal{W}; \mathcal{K} \rrbracket$, respectively.
 1060 We have by the Cauchy-Schwarz inequality that

$$\left(\sum_{d=1}^{D_{\mathcal{K}}} |w_d^2 - v_d^2| \right)^2 = \left(\sum_{d=1}^{D_{\mathcal{K}}} |w_d - v_d| \cdot |w_d + v_d| \right)^2 \leq \left(\sum_{d=1}^{D_{\mathcal{K}}} (w_d - v_d)^2 \right) \left(\sum_{d=1}^{D_{\mathcal{K}}} (w_d + v_d)^2 \right).$$

1061 Now the first term is upper bounded by $\|\mathcal{V} - \mathcal{W}\|^2$, and for the second we have

$$\sum_{d=1}^{D_{\mathcal{K}}} (w_d + v_d)^2 = \sum_{d=1}^{D_{\mathcal{K}}} w_d^2 + \sum_{d=1}^{D_{\mathcal{K}}} v_d^2 + 2v_d w_d = 2 + 2 \sum_{d=1}^{D_{\mathcal{K}}} v_d w_d \leq 4,$$

1062 where we use the fact that $\|\mathcal{V}\| = \|\mathcal{W}\| = 1$, and again Cauchy-Schwarz. Overall we have:

$$\sum_{d=1}^{D_{\mathcal{K}}} |w_d^2 - v_d^2| \leq 2\|\mathcal{V} - \mathcal{W}\|.$$

1063 Note that the left hand side of the inequality above equals twice the total variation distance between
 1064 the distributions defined by $\{w_d^2\}_{d=1}^{D_{\mathcal{K}}}$ and $\{v_d^2\}_{d=1}^{D_{\mathcal{K}}}$. Therefore by Lemma 2 we have:

$$|QE(\mathcal{V}; \mathcal{K}) - QE(\mathcal{W}; \mathcal{K})| = |H(\{w_d^2\}) - H(\{v_d^2\})| \leq \|\mathcal{V} - \mathcal{W}\| \cdot \ln(D_{\mathcal{K}}) + H_b(\|\mathcal{V} - \mathcal{W}\|).$$

1065 \square

1066 **Lemma 9.** Let $\mathcal{P} = \{p(x)\}_{x \in [S]}$, where $S \in \mathbb{N}$, be a probability distribution, and denote its entropy
 1067 by $H(\mathcal{P}) := \mathbb{E}_{x \sim \mathcal{P}}[\ln(1/p(x))]$. Then, for any $0 < a < 1$, there exists a subset $T \subseteq [S]$ such that
 1068 $Pr_{x \sim \mathcal{P}}(T^c) \leq a$ and $|T| \leq e^{\frac{H(\mathcal{P})}{a}}$.

1069 *Proof.* By Markov's inequality we have for any $0 < a < 1$:

$$Pr_{x \sim \mathcal{P}} \left(\left\{ x : e^{-\frac{H(\mathcal{P})}{a}} \geq p(x) \right\} \right) = Pr_{x \sim \mathcal{P}} \left(\left\{ x : \ln \left(\frac{1}{p(x)} \right) \geq \frac{H(\mathcal{P})}{a} \right\} \right) \leq a.$$

1070 Let $T := \{x : e^{-\frac{H(\mathcal{P})}{a}} \leq p(x)\} \subseteq [S]$. Note that

$$e^{-\frac{H(\mathcal{P})}{a}} |T| \leq \sum_{x \in T} p(x) \leq \sum_{x \in [S]} p(x) = 1,$$

1071 and so $|T| \leq e^{\frac{H(\mathcal{P})}{a}}$ and $Pr_{x \sim \mathcal{P}}(T^c) \leq a$, as required. \square

1072 **K.2 Proof of Theorem 1**

1073 If $\mathcal{A} = 0$ the theorem is trivial, since then $QE(\mathcal{A}; \mathcal{K}) = 0$ for all $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$, so we can assume
 1074 $\mathcal{A} \neq 0$. We have:

$$\begin{aligned} \left\| \frac{\mathcal{W}_{\text{TN}}}{\|\mathcal{W}_{\text{TN}}\|} - \frac{\mathcal{A}}{\|\mathcal{A}\|} \right\| &= \frac{1}{\|\mathcal{A}\|} \left\| \frac{\|\mathcal{A}\|}{\|\mathcal{W}_{\text{TN}}\|} \cdot \mathcal{W}_{\text{TN}} - \mathcal{A} \right\| \\ &\leq \frac{1}{\|\mathcal{A}\|} \left(\left| \frac{\|\mathcal{A}\|}{\|\mathcal{W}_{\text{TN}}\|} - 1 \right| \cdot \|\mathcal{W}_{\text{TN}}\| + \|\mathcal{W}_{\text{TN}} - \mathcal{A}\| \right) \\ &= \frac{1}{\|\mathcal{A}\|} (\|\mathcal{A}\| - \|\mathcal{W}_{\text{TN}}\| + \|\mathcal{W}_{\text{TN}} - \mathcal{A}\|) \\ &\leq \frac{2\epsilon}{\|\mathcal{A}\|}. \end{aligned}$$

1075 Now, let $\hat{\mathcal{A}} := \frac{\mathcal{A}}{\|\mathcal{A}\|}$ and $\hat{\mathcal{W}}_{\text{TN}} = \frac{\mathcal{W}_{\text{TN}}}{\|\mathcal{W}_{\text{TN}}\|}$ be normalized versions of \mathcal{A} and \mathcal{W}_{TN} , respectively, and let
 1076 $c = \frac{2\epsilon}{\|\mathcal{A}\|}$. Note that $c < \frac{2\|\mathcal{A}\|}{4\|\mathcal{A}\|} = \frac{1}{2}$, and therefore by Lemma 8 we have:

$$\begin{aligned} |QE(\mathcal{A}; \mathcal{K}) - QE(\mathcal{W}_{\text{TN}}; \mathcal{K})| &= \left| QE(\hat{\mathcal{A}}; \mathcal{K}) - QE(\hat{\mathcal{W}}_{\text{TN}}; \mathcal{K}) \right| \\ &\leq c \cdot \ln(D_{\mathcal{K}}) + H_b(c). \end{aligned}$$

1077 By Lemma 4 we have that

$$QE(\hat{\mathcal{W}}_{\text{TN}}; \mathcal{K}) \leq \ln(\text{rank}(\llbracket \mathcal{W}_{\text{TN}}; \mathcal{K} \rrbracket)) \leq \ln(R),$$

1078 and therefore

$$QE(\hat{\mathcal{A}}; \mathcal{K}) \leq \ln(R) + c \ln(D_{\mathcal{K}}) + H_b(c).$$

1079 Substituting $c = \frac{2\epsilon}{\|\mathcal{A}\|}$ and invoking the elementary inequality $H_b(x) \leq 2\sqrt{x}$ we obtain

$$QE(\mathcal{A}; \mathcal{K}) \leq \ln(R) + \frac{2\epsilon}{\|\mathcal{A}\|} \cdot \ln(D_{\mathcal{K}}) + 2\sqrt{\frac{2\epsilon}{\|\mathcal{A}\|}},$$

1080 as required. Lastly, the existence of a tensor $\mathcal{A}' \in \mathbb{R}^{D_1 \times \dots \times D_N}$ such that $QE(\mathcal{A}'; \mathcal{K}) \geq$
 1081 $\min\{|\mathcal{K}|, |\mathcal{K}^c|\} \cdot \ln(\min_{n \in [N]} D_n)$ for all $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$ follows from the construction given in
 1082 Theorem 2 of [17]. \square

1083 **K.3 Proof of Theorem 2**

1084 Consider, for each canonical partition $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$, the distribution

$$\mathcal{P}_{\mathcal{K}} = \left\{ p_{\mathcal{K}}(i) := \frac{\sigma_{\mathcal{K},i}^2}{\|\mathcal{A}\|^2} \right\}_{i \in [D_{\mathcal{K}}]},$$

1085 where $\sigma_{\mathcal{K},1} \geq \sigma_{\mathcal{K},2} \geq \dots \geq \sigma_{\mathcal{K},D_{\mathcal{K}}}$ are the singular values of $\llbracket \mathcal{A}; \mathcal{K} \rrbracket$ (note that $\frac{1}{\|\mathcal{A}\|^2} \sum_j \sigma_{\mathcal{K},j}^2 =$
 1086 $\frac{\|\mathcal{A}\|^2}{\|\mathcal{A}\|^2} = 1$ so $\mathcal{P}_{\mathcal{K}}$ is indeed a probability distribution). Denoting by $H(\mathcal{P}_{\mathcal{K}}) := \mathbb{E}_{i \sim \mathcal{P}_{\mathcal{K}}}[\ln(1/p_{\mathcal{K}}(i))]$
 1087 the entropy of $\mathcal{P}_{\mathcal{K}}$, by assumption:

$$QE(\mathcal{A}; \mathcal{K}) = H(\mathcal{P}_{\mathcal{K}}) \leq \frac{\epsilon^2}{\|\mathcal{A}\|^2(2N-3)} \ln(R),$$

1088 for all $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$. Thus, taking $a = \frac{\epsilon^2}{\|\mathcal{A}\|^2(2N-3)}$ we obtain by Lemma 9 that there exists a subset
 1089 $T_{\mathcal{K}} \subseteq [D_{\mathcal{K}}]$ such that

$$\mathcal{P}_{\mathcal{K}}(T_{\mathcal{K}}^c) \leq \frac{\epsilon^2}{(2N-3)\|\mathcal{A}\|^2},$$

1090 and $|T_{\mathcal{K}}| \leq e^{\frac{H(\mathcal{P}_{\mathcal{K}})}{a}} = e^{\ln(R)} = R$. Note that

$$\mathcal{P}_{\mathcal{K}}(T_{\mathcal{K}}) \leq \sum_{i=1}^R \frac{\sigma_i^2}{\|\mathcal{A}\|^2}.$$

1091 Since this holds for any subset of cardinality at most R . Taking complements we obtain

$$\sum_{i=R+1}^{D_{\mathcal{K}}} \frac{\sigma_i^2}{\|\mathcal{A}\|^2} \leq \mathcal{P}_{\mathcal{K}}(T_{\mathcal{K}}^{\epsilon}),$$

1092 so

$$\sqrt{\sum_{i=R+1}^{D_{\mathcal{K}}} \sigma_{\mathcal{K},i}^2} \leq \frac{\epsilon}{\sqrt{(2N-3)}}.$$

1093 We can now invoke Lemma 6, which implies that there exists some $\mathcal{W}_{\text{TN}} \in \mathbb{R}^{D_1 \times \dots \times D_N}$ generated
1094 by the locally connected tensor network satisfying:

$$\|\mathcal{W}_{\text{TN}} - \mathcal{A}\| \leq \epsilon.$$

1095

□

1096 **K.4 Proof of Proposition 1**

1097 Recall that by Lemma 8 for any pair of unit norm tensors \mathcal{V}, \mathcal{W} such that $\|\mathcal{V} - \mathcal{W}\| < 0.5$

$$|QE(\mathcal{V}; \mathcal{K}) - QE(\mathcal{W}; \mathcal{K})| \leq \|\mathcal{V} - \mathcal{W}\| \ln(\mathcal{D}_{\mathcal{K}}) + H_b(\|\mathcal{V} - \mathcal{W}\|) \leq 2 \ln(\mathcal{D}_{\mathcal{K}}) \sqrt{\|\mathcal{V} - \mathcal{W}\|},$$

1098 where we used the elementary inequality $H_b(x) \leq 2\sqrt{x}$. So it suffices to bound

$$\left\| \frac{\mathcal{D}_{\text{emp}}}{\|\mathcal{D}_{\text{pop}}\|} - \frac{\mathcal{D}_{\text{emp}}}{\|\mathcal{D}_{\text{emp}}\|} \right\|.$$

1099 We have the identity

$$\left\| \frac{\mathcal{D}_{\text{pop}}}{\|\mathcal{D}_{\text{pop}}\|} - \frac{\mathcal{D}_{\text{emp}}}{\|\mathcal{D}_{\text{emp}}\|} \right\| = \left\| \frac{\mathcal{D}_{\text{pop}} \|\mathcal{D}_{\text{emp}}\| - \mathcal{D}_{\text{emp}} \|\mathcal{D}_{\text{pop}}\|}{\|\mathcal{D}_{\text{pop}}\| \|\mathcal{D}_{\text{emp}}\|} \right\| =$$

1100

$$\left\| \frac{\mathcal{D}_{\text{pop}} \|\mathcal{D}_{\text{emp}}\| - \|\mathcal{D}_{\text{emp}}\| \|\mathcal{D}_{\text{emp}} + \mathcal{D}_{\text{emp}}\| \|\mathcal{D}_{\text{emp}}\| - \mathcal{D}_{\text{emp}} \|\mathcal{D}_{\text{pop}}\|}{\|\mathcal{D}_{\text{pop}}\| \|\mathcal{D}_{\text{emp}}\|} \right\|$$

1101 by the triangle inequality the above is bounded by

$$\frac{\|\mathcal{D}_{\text{pop}} - \mathcal{D}_{\text{emp}}\|}{\|\mathcal{D}_{\text{pop}}\|} + \frac{\|\|\mathcal{D}_{\text{pop}}\| - \|\mathcal{D}_{\text{emp}}\|\|}{\|\mathcal{D}_{\text{pop}}\|}$$

1102 For $m \in [M]$, let $\mathcal{X}^{(m)} = y^{(m)} \cdot \otimes_{n \in [N]} \mathbf{x}^{(n,m)} - \mathcal{D}_{\text{pop}}$. These are i.i.d random variables with
1103 $\mathbb{E}[\mathcal{X}^{(m)}] = 0$ and $\|\mathcal{X}^{(m)}\| \leq 2$ for all $m \in [M]$. Note that

$$\left\| \frac{1}{M} \sum_{m=1}^M \mathcal{X}^{(m)} \right\| = \|\mathcal{D}_{\text{emp}} - \mathcal{D}_{\text{pop}}\|,$$

1104 so by Lemma 2 with $c = 2, t = \frac{\|\mathcal{D}_{\text{pop}}\| \gamma}{\ln(\mathcal{D}_{\mathcal{K}})}$, assuming $M \geq \frac{128 \ln(\frac{2}{\delta}) (\ln(\mathcal{D}_{\mathcal{K}}))^4}{\|\mathcal{D}_{\text{pop}}\|^2 \gamma^4}$ we have with probability
1105 at least $1 - \delta$

$$\|\|\mathcal{D}_{\text{pop}}\| - \|\mathcal{D}_{\text{emp}}\|\| \leq \|\mathcal{D}_{\text{pop}} - \mathcal{D}_{\text{emp}}\| \leq \frac{\|\mathcal{D}_{\text{pop}}\| \gamma^2}{8 \ln(\mathcal{D}_{\mathcal{K}})^2},$$

1106 and therefore

$$\frac{\|\mathcal{D}_{\text{pop}} - \mathcal{D}_{\text{emp}}\|}{\|\mathcal{D}_{\text{pop}}\|} + \frac{\|\|\mathcal{D}_{\text{pop}}\| - \|\mathcal{D}_{\text{emp}}\|\|}{\|\mathcal{D}_{\text{pop}}\|} \leq \frac{\gamma^2}{4 \ln(\mathcal{D}_{\mathcal{K}})^2},$$

1107 and so by Lemma 8

$$|QE(\mathcal{D}_{\text{emp}}; \mathcal{K}) - QE(\mathcal{D}_{\text{pop}}; \mathcal{K})| \leq \gamma.$$

1108

□

1109 **K.5 Proof of Corollary 1**

1110 Notice that the entanglements of a tensor are invariant to multiplication by a constant. In particular,
 1111 $QE(\mathcal{D}_{\text{pop}}; \mathcal{K}) = QE(\mathcal{D}_{\text{pop}}/\|\mathcal{D}_{\text{pop}}\|; \mathcal{K})$ for any $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$. Hence, if there exists a canonical
 1112 partition $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$ under which $QE(\mathcal{D}_{\text{pop}}; \mathcal{K}) > \ln(R) + 2\epsilon \cdot \ln(D_{\mathcal{K}}) + 2\sqrt{2}\epsilon$, then Theorem 1
 1113 implies that $\min_{\mathcal{W}_{\text{TN}}} \|\mathcal{W}_{\text{TN}} - \mathcal{D}_{\text{pop}}/\|\mathcal{D}_{\text{pop}}\|\| > \epsilon$. Now, for any non-zero $\mathcal{W} \in \mathbb{R}^{D_1 \times \dots \times D_N}$ gener-
 1114 ated by the locally connected tensor network, one can also represent $\mathcal{W}/\|\mathcal{W}\|$ by multiplying any of
 1115 the tensors constituting the tensor network by $1/\|\mathcal{W}\|$ (contraction is a multilinear operation). Thus:

$$\text{SubOpt} := \min_{\mathcal{W}_{\text{TN}}} \left\| \frac{\mathcal{W}_{\text{TN}}}{\|\mathcal{W}_{\text{TN}}\|} - \frac{\mathcal{D}_{\text{pop}}}{\|\mathcal{D}_{\text{pop}}\|} \right\| \geq \min_{\mathcal{W}_{\text{TN}}} \left\| \mathcal{W}_{\text{TN}} - \frac{\mathcal{D}_{\text{pop}}}{\|\mathcal{D}_{\text{pop}}\|} \right\| > \epsilon,$$

1116 which concludes the first part of the claim, *i.e.* the necessary condition for low suboptimality in
 1117 achievable accuracy.

1118 For the sufficient condition, if for all $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N$ it holds that $QE(\mathcal{D}_{\text{pop}}; \mathcal{K}) \leq \frac{\epsilon^2}{8N-12} \cdot \ln(R)$,
 1119 then by Theorem 2 there exists an assignment for the locally connected tensor network such that
 1120 $\|\mathcal{W}_{\text{TN}} - \mathcal{D}_{\text{pop}}/\|\mathcal{D}_{\text{pop}}\|\| \leq \epsilon/2$. From the triangle inequality we obtain:

$$\left\| \frac{\mathcal{W}_{\text{TN}}}{\|\mathcal{W}_{\text{TN}}\|} - \frac{\mathcal{D}_{\text{pop}}}{\|\mathcal{D}_{\text{pop}}\|} \right\| \leq \left\| \mathcal{W}_{\text{TN}} - \frac{\mathcal{D}_{\text{pop}}}{\|\mathcal{D}_{\text{pop}}\|} \right\| + \left\| \mathcal{W}_{\text{TN}} - \frac{\mathcal{W}_{\text{TN}}}{\|\mathcal{W}_{\text{TN}}\|} \right\| \leq \frac{\epsilon}{2} + \left\| \mathcal{W}_{\text{TN}} - \frac{\mathcal{W}_{\text{TN}}}{\|\mathcal{W}_{\text{TN}}\|} \right\|. \quad (10)$$

1121 Since $\|\mathcal{W}_{\text{TN}} - \mathcal{D}_{\text{pop}}/\|\mathcal{D}_{\text{pop}}\|\| \leq \epsilon/2$ it holds that $\|\mathcal{W}_{\text{TN}}\| \leq 1 + \epsilon/2$. Combined with the fact that
 1122 $\left\| \mathcal{W}_{\text{TN}} - \frac{\mathcal{W}_{\text{TN}}}{\|\mathcal{W}_{\text{TN}}\|} \right\| = \|\mathcal{W}_{\text{TN}}\| - 1$, we get that $\left\| \mathcal{W}_{\text{TN}} - \frac{\mathcal{W}_{\text{TN}}}{\|\mathcal{W}_{\text{TN}}\|} \right\| \leq \epsilon/2$. Plugging this into Equa-
 1123 tion (10) yields:

$$\left\| \frac{\mathcal{W}_{\text{TN}}}{\|\mathcal{W}_{\text{TN}}\|} - \frac{\mathcal{D}_{\text{pop}}}{\|\mathcal{D}_{\text{pop}}\|} \right\| \leq \epsilon,$$

1124 and so $\text{SubOpt} := \min_{\mathcal{W}_{\text{TN}}} \left\| \frac{\mathcal{W}_{\text{TN}}}{\|\mathcal{W}_{\text{TN}}\|} - \frac{\mathcal{D}_{\text{pop}}}{\|\mathcal{D}_{\text{pop}}\|} \right\| \leq \epsilon$. □

1125 **K.6 Proof of Theorem 3**

1126 If $\mathcal{A} = 0$ the theorem is trivial, since then $QE(\mathcal{A}; \mu(\mathcal{K})) = 0$ for all $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N^P$, so we can
 1127 assume $\mathcal{A} \neq 0$. We have:

$$\begin{aligned} \left\| \frac{\mathcal{W}_{\text{TN}}^P}{\|\mathcal{W}_{\text{TN}}^P\|} - \frac{\mathcal{A}}{\|\mathcal{A}\|} \right\| &= \frac{1}{\|\mathcal{A}\|} \left\| \frac{\|\mathcal{A}\|}{\|\mathcal{W}_{\text{TN}}^P\|} \cdot \mathcal{W}_{\text{TN}}^P - \mathcal{A} \right\| \\ &\leq \frac{1}{\|\mathcal{A}\|} \left(\left| \frac{\|\mathcal{A}\|}{\|\mathcal{W}_{\text{TN}}^P\|} - 1 \right| \cdot \|\mathcal{W}_{\text{TN}}^P\| + \|\mathcal{W}_{\text{TN}}^P - \mathcal{A}\| \right) \\ &= \frac{1}{\|\mathcal{A}\|} (\|\mathcal{A}\| - \|\mathcal{W}_{\text{TN}}^P\| + \|\mathcal{W}_{\text{TN}}^P - \mathcal{A}\|) \\ &\leq \frac{2\epsilon}{\|\mathcal{A}\|}. \end{aligned}$$

1128 Now, let $\hat{\mathcal{A}} := \frac{\mathcal{A}}{\|\mathcal{A}\|}$ and $\hat{\mathcal{W}}_{\text{TN}}^P = \frac{\mathcal{W}_{\text{TN}}^P}{\|\mathcal{W}_{\text{TN}}^P\|}$ be normalized versions of \mathcal{A} and $\mathcal{W}_{\text{TN}}^P$, respectively, and let
 1129 $c = \frac{2\epsilon}{\|\mathcal{A}\|}$. Note that $c < \frac{2\|\mathcal{A}\|}{4} \frac{1}{\|\mathcal{A}\|} = \frac{1}{2}$, and therefore by Lemma 8 we have:

$$\begin{aligned} |QE(\mathcal{A}; \mu(\mathcal{K})) - QE(\mathcal{W}_{\text{TN}}^P; \mu(\mathcal{K}))| &= \left| QE(\hat{\mathcal{A}}; \mu(\mathcal{K})) - QE(\hat{\mathcal{W}}_{\text{TN}}^P; \mu(\mathcal{K})) \right| \\ &\leq c \cdot \ln(D_{\mu(\mathcal{K})}) + H_b(c). \end{aligned}$$

1130 By Lemma 5 we have that

$$QE(\hat{\mathcal{W}}_{\text{TN}}^P; \mu(\mathcal{K})) \leq \ln(\text{rank}(\llbracket \mathcal{W}_{\text{TN}}^P; \mu(\mathcal{K}) \rrbracket)) \leq \ln(R),$$

1131 and therefore

$$QE(\hat{\mathcal{A}}; \mu(\mathcal{K})) \leq \ln(R) + c \ln(D_{\mu(\mathcal{K})}) + H_b(c).$$

1132 Substituting $c = \frac{2\epsilon}{\|\mathcal{A}\|}$ and invoking the elementary inequality $H_b(x) \leq 2\sqrt{x}$ we obtain

$$QE(\mathcal{A}; \mu(\mathcal{K})) \leq \ln(R) + \frac{2\epsilon}{\|\mathcal{A}\|} \cdot \ln(D_{\mu(\mathcal{K})}) + 2\sqrt{\frac{2\epsilon}{\|\mathcal{A}\|}},$$

1133 as required. Equation (6) follows from the construction given in Theorem 2 of [17].

1134

□

1135 **K.7 Proof of Theorem 4**

1136 Let \mathcal{C}_{N^P} be the one-dimensional canonical partitions of $[N^P]$ (Definition 2). Note that $\mu(\mathcal{C}_N^P) :=$
 1137 $\{\mu(\mathcal{K}) : \mathcal{K} \in \mathcal{C}_N^P\} \subseteq \mathcal{C}_{N^P}$. For an assignment $(n_{\mathcal{K}})_{\mathcal{K} \in \mathcal{C}_{N^P}} \in \mathbb{N}^{\mathcal{C}_{N^P}}$ of integers to one-dimensional
 1138 canonical partitions $\mathcal{K} \in \mathcal{C}_{N^P}$, we consider the set of tensors whose matricization with respect to
 1139 each $\mathcal{K} \in \mathcal{C}_{N^P}$ has rank at most $n_{\mathcal{K}}$. This set is also known in the literature as the set of tensors with
 1140 *Hierarchical Tucker (HT) rank* at most $(n_{\mathcal{K}})_{\mathcal{K} \in \mathcal{C}_{N^P}}$ (cf. [25]). Accordingly, we denote it by:

$$HT((n_{\mathcal{K}})_{\mathcal{K} \in \mathcal{C}_{N^P}}) := \{\mathcal{V} \in \mathbb{R}^{D_1 \times \dots \times D_{N^P}} : \forall \mathcal{K} \in \mathcal{C}_{N^P}, \text{rank}(\llbracket \mathcal{V}; \mathcal{K} \rrbracket) \leq n_{\mathcal{K}}\}.$$

1141 Now, define $(n_{\mathcal{K}}^*)_{\mathcal{K} \in \mathcal{C}_{N^P}} \in \mathbb{N}^{\mathcal{C}_{N^P}}$ by:

$$\forall \mathcal{K} \in \mathcal{C}_{N^P} : n_{\mathcal{K}}^* = \begin{cases} R & \text{if } \mu^{-1}(\mathcal{K}) \in \mathcal{C}_N^P \\ \mathcal{D}_{\mathcal{K}} & \text{if } \mu^{-1}(\mathcal{K}) \notin \mathcal{C}_N^P \end{cases},$$

1142 where $\mathcal{D}_{\mathcal{K}} := \min\{\prod_{n \in \mathcal{K}} D_n, \prod_{n \in \mathcal{K}^c} D_n\}$. We show that for any tensor \mathcal{A} that satisfies for all
 1143 $\mathcal{K} \in \mathcal{C}_N^P$:

$$QE(\mathcal{A}; \mu(\mathcal{K})) \leq \frac{\epsilon^2}{(2N^P - 3)\|\mathcal{A}\|^2} \cdot \ln(R),$$

1144 there exists a tensor $\mathcal{V} \in HT((n_{\mathcal{K}}^*)_{\mathcal{K} \in \mathcal{C}_{N^P}})$ such that $\|\mathcal{A} - \mathcal{V}\| \leq \epsilon$. Consider, for each canonical
 1145 partition $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N^P$, the distribution

$$\mathcal{P}_{\mathcal{K}} = \left\{ p_{\mathcal{K}}(i) := \frac{\sigma_{\mathcal{K},i}^2}{\|\mathcal{A}\|^2} \right\}_{i \in [D_{\mathcal{K}}]},$$

1146 where $\sigma_{\mathcal{K},1} \geq \sigma_{\mathcal{K},2} \geq \dots \geq \sigma_{\mathcal{K},D_{\mathcal{K}}}$ are the singular values of $\llbracket \mathcal{A}; \mu(\mathcal{K}) \rrbracket$ (note that $\frac{1}{\|\mathcal{A}\|^2} \sum_j \sigma_{\mathcal{K},j}^2 =$
 1147 $\frac{\|\mathcal{A}\|^2}{\|\mathcal{A}\|^2} = 1$ so $\mathcal{P}_{\mathcal{K}}$ is indeed a probability distribution). Denoting by $H(\mathcal{P}_{\mathcal{K}}) := \mathbb{E}_{i \sim \mathcal{P}_{\mathcal{K}}}[\ln(1/p_{\mathcal{K}}(i))]$
 1148 the entropy of $\mathcal{P}_{\mathcal{K}}$, by assumption:

$$QE(\mathcal{A}; \mu(\mathcal{K})) = H(\mathcal{P}_{\mathcal{K}}) \leq \frac{\epsilon^2}{\|\mathcal{A}\|^2(2N^P - 3)} \ln(R),$$

1149 for all $(\mathcal{K}, \mathcal{K}^c) \in \mathcal{C}_N^P$. Thus, taking $a = \frac{\epsilon^2}{\|\mathcal{A}\|^2(2N^P - 3)}$ we get by Lemma 9 that there exists a subset
 1150 $T_{\mathcal{K}} \subseteq [D_{\mathcal{K}}]$ such that

$$\mathcal{P}_{\mathcal{K}}(T_{\mathcal{K}}^c) \leq \frac{\epsilon^2}{(2N^P - 3)\|\mathcal{A}\|^2},$$

1151 and $|T_{\mathcal{K}}| \leq e^{\frac{H(\mathcal{P}_{\mathcal{K}})}{a}} = e^{\ln(R)} = R$. Note that

$$\mathcal{P}_{\mathcal{K}}(T_{\mathcal{K}}) \leq \sum_{i=1}^R \frac{\sigma_i^2}{\|\mathcal{A}\|^2}.$$

1152 Since this holds for any subset of cardinality at most R . Taking complements we obtain

$$\sum_{i=R+1}^{D_{\mathcal{K}}} \frac{\sigma_i^2}{\|\mathcal{A}\|^2} \leq \mathcal{P}_{\mathcal{K}}(T_{\mathcal{K}}^c),$$

1153 so

$$\sqrt{\sum_{i=R+1}^{D_{\mathcal{K}}} \sigma_{\mathcal{K},i}^2} \leq \frac{\epsilon}{\sqrt{(2N^P - 3)}}.$$

1154 We can now invoke Lemma 6 (note that if $\mu^{-1}(\mathcal{K}) \notin \mathcal{C}_N^P$, then the requirements of Lemma 6 are
 1155 trivially fulfilled with respect to the partition $(\mathcal{K}, \mathcal{K}^c)$ since $n_{\mathcal{K}}^* = D_{\mathcal{K}}$), which implies that there
 1156 exists some $\mathcal{W} \in HT((n_{\mathcal{K}}^*)_{\mathcal{K} \in \mathcal{C}_{N^P}})$ satisfying:

$$\|\mathcal{W} - \mathcal{A}\| \leq \epsilon,$$

1157 as required.

1158 The proof concludes by establishing that for any tensor $\mathcal{W} \in HT((n_{\mathcal{K}}^*)_{\mathcal{K} \in \mathcal{C}_{N^P}})$, there exists
1159 assignment for the tensors constituting the P -dimensional locally connected tensor network (defined
1160 in Figure 6) such that it generates \mathcal{W} .

1161 To see why this is the case, note that by Lemma 7 any tensor $\mathcal{W} \in HT((n_{\mathcal{K}}^*)_{\mathcal{K} \in \mathcal{C}_{N^P}})$ can be repre-
1162 sented by a (one-dimensional) locally connected tensor network with varying widths $(n_{\mathcal{K}}^*)_{\mathcal{K} \in \mathcal{C}_{N^P}}$,
1163 *i.e.* a tensor network conforming to a perfect binary tree graph in which the lengths of inner axes are
1164 as follows: an axis corresponding to an edge that connects a node with descendant leaves indexed by
1165 \mathcal{K} to its parent is assigned the length $n_{\mathcal{K}}^*$. We can obtain an equivalent representation of any such
1166 tensor as a P -dimensional locally connected tensor network (described in Appendix I.1.1) via the
1167 following procedure. For each node at level $l \in \{0, P, 2P, \dots, (L-1)P\}$ of the tree (recall $N = 2^L$
1168 for $L \in \mathbb{N}$), contract it with its descendant nodes at levels $\{l+1, \dots, l+(P-1)\}$.⁹ This results in
1169 a new tensor network whose underlying graph is a perfect 2^P -ary tree and the remaining edges all
1170 correspond to inner axes of lengths equal to $n_{\mathcal{K}}^* = R$ for $\mathcal{K} \in \mathcal{C}_N^P$, *i.e.* in a representation of \mathcal{W} as a
1171 P -dimensional locally connected tensor network. \square

⁹For a concrete example, let $N = 2^L = 4$ and $P = 2$ (*i.e.* $L = 2$). In this case, the perfect binary tree underlying the one-dimensional locally connected tensor network of varying widths is of height $L \cdot P = 4$ and has $N^P = 16$ leaves. It is converted into a perfect 4-ary tree tensor network of height $L = 2$ by contracting the root with its two children and the nodes at level two with their children.