Chapter 3

Advanced Tensor Network Manipulations

A. Special Cases

Last day someone asked about if the eigenvalues of the transfer matrix were always real, and if the transfer matrix is always diagonalizable. The answer to both those questions is "no." I'll just give you the examples, and let you figure out (1) why and (2) how expectation values look. It is a good exercise. Both of these examples are what is referred to as "non-invective" – meaning that the boundary conditions matter.

The "single-particle" wavefunction has a non-diagonalizable transfer matrix:

$$|\psi\rangle = \cdots \begin{pmatrix} |0\rangle & \frac{1}{\sqrt{N}}|1\rangle \\ & |0\rangle \end{pmatrix} \begin{pmatrix} |0\rangle & \frac{1}{\sqrt{N}}|1\rangle \\ & |0\rangle \end{pmatrix} \begin{pmatrix} |0\rangle & \frac{1}{\sqrt{N}}|1\rangle \\ & |0\rangle \end{pmatrix} \cdots \tag{3.1}$$

$$= \frac{1}{\sqrt{N}} \left(\cdots + \cdots 100 \cdots + \cdots 010 \cdots + \cdots 001 \cdots + \cdots \right). \tag{3.2}$$

That is, it is the equal sum over all configurations with a single $|1|\rangle$.

The 3-state chiral Schrodinger Cat has complex eigenvalues:

$$|\psi\rangle = \cdots \begin{pmatrix} & |A\rangle & \\ & & |B\rangle \\ |C\rangle & & \end{pmatrix} \begin{pmatrix} & |A\rangle & \\ & & |B\rangle \\ |C\rangle & & \end{pmatrix} \begin{pmatrix} & |A\rangle & \\ & & |B\rangle \\ |C\rangle & & \end{pmatrix} \cdots \tag{3.3}$$

$$= \cdots ABCABCABC \cdots + \cdots BCABCABCA \cdots + \cdots CABCABCAB \cdots . \tag{3.4}$$

B. Programming

Do the third intro to tensors notebook.

C. Contracting Tensors

As we saw last day, one can use tensor networks to describe quantum mechanical wavefunctions. A common operation one needs to do with tensor networks is contract them. In the examples we looked at so far, we did this in an ad-hoc way. We are now going to explore the computational cost of contracting tensors, and will see that the order in which one contracts a tensor network can make a big difference on the cost.

How many operations are needed to take a dot-product? Multiply a matrix and a vector? Multiply two matrices? Multiply two tensors?

Lets begin with the simplest contraction, a dot product between two vectors:

$$\vec{v} \cdot \vec{w} = \tag{3.5}$$

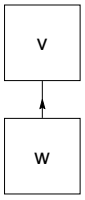

If the index being contracted can take on d separate values, then we need to perform d independent multiplication, and then add these numbers. This can be parallelized, but if we think about serial operations, this is of order d operations. Before the operation we are storing $2d$ numbers, after it we are storing 1 number.

Now consider multiplying a matrix times a vector. If the matrix is $d_1 \times d_2$, and the vector is length d_2, we need to do $d_1 \times d_2$ multiplications, and a bunch of additions. Before the operation we are storing $(d_1 + 1) \times d_2$ numbers, and afterwards, just d_2.

Matrix multiplications: A $d_1 \times d_2$ matrix multiplied by a $d_2 \times d_3$ requires $d_1 \times d_2 \times d_3$ operations.

Tensor multiplication: A $d_1 \times d_2 \times \cdots d_n$ tensor where we contract the n'th index with at $d_n \times d_{n+1} \times \cdots d_{n+m}$ tensor requires $d_1 \times d_2 \times \cdot d_{n+m}$ operations, and result in $d_1 \times d_2 \times \cdot d_{n+m}/d_n$ numbers to be stored. Tensor multiplication for high rank tensors can become expensive – moreover, the number of numbers that you have to store goes up as you multiply the tensors together. In some sense this is why tensor networks are useful for encoding quantum states: they are some sort of compression method.

A consequence of this logic is that the order that you contract a tensor network changes the cost. Consider one of the contractions from last day:

$$\langle\psi|\psi\rangle = \tag{3.6}$$

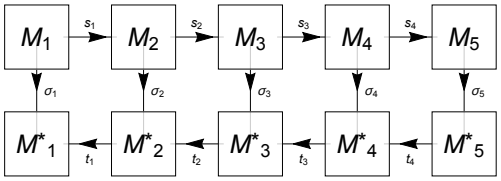

Suppose all of the s and t's take on $\chi = 3$ possible values, and all the σ's take on $d = 2$ possible values, what is the most efficient way to contract this network?

The cheapest one to contract is σ_1 – which takes $3 \times 3 \times 2$ operations. After the contraction, the network looks like

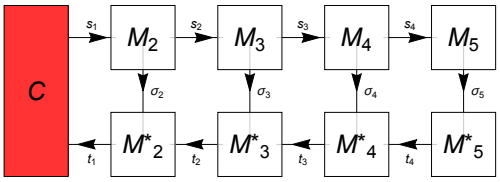

Next we do s_1, which takes $3 \times 3 \times 2$ operation, and yields

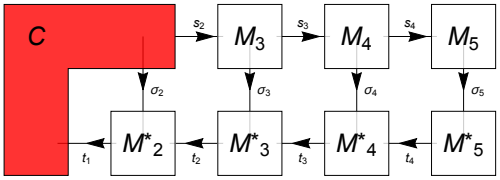

We then contract s_1 and σ_2. These can be done simultaneously, or one at a time. If we do then simultaneously, there are $3 \times 3 \times 3 \times 2$ operations. If we do them one at a time, the first one takes $3 \times 3 \times 3 \times 2$ operations, and the second $3 \times 3 \times 2$ operations. Clearly there is no gain from doing them one at a time.

In the end it takes of order $N \times \chi^3 \times d$ operations to take the overlap between two matrix product states. This is pretty good, considering the Hilbert space is exponential in N.

The most natural generalization of a matrix product state to two dimensions is a "PEPS" – which is a tensor network which look like

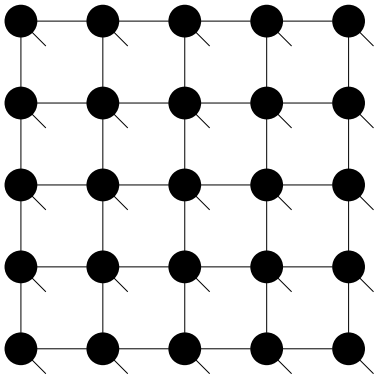

The dots represent rank 5 tensors, and the dangling bonds are the physical spin indices. Exactly contracting a tensor network like this is extremely expensive. Imagine we just need to contract something like

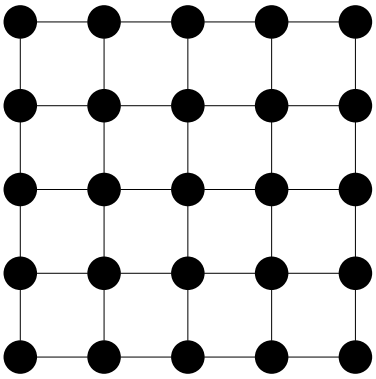

Estimate the number of operations needed to contract this network.

The problem is that at some point one needs to have a very high rank tensor, such as

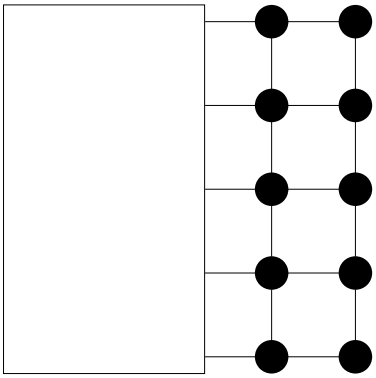

Just storing this tensor, let alone operating with it, requires χ^L numbers, where L is the linear size of the system. Contracting such a network requires a number of operations which scales as the exponential of the square root of the total number of tensors. One important area is the development of algorithms to efficiently approximate such contractions. Another is finding tensor networks which are easy to contract, but do a good job of representing quantum mechanical states

D. Singular Value Decomposition

In addition to the operation of contraction, it is important to have the opposite operation: breaking a high ranking tensor into several lower ranking ones, for example,

The tool we will use for this is the singular value decomposition (SVD). This is an algorithm which lets you write a $n \times m$ matrix as the product

$$M = U\Lambda V \tag{3.7}$$

where U is $n \times d$ matrix, Λ is a diagonal $d \times d$ matrix, and V is $d \times m$, with $d \leq \min(m, n)$. The entries in

Λ are the "singular values," the d rows of V are orthonormal, and the d columns of U are orthonormal:

$$VV^\dagger = I \tag{3.8}$$

$$U^\dagger U = I \tag{3.9}$$

Graphically these relations can be written:

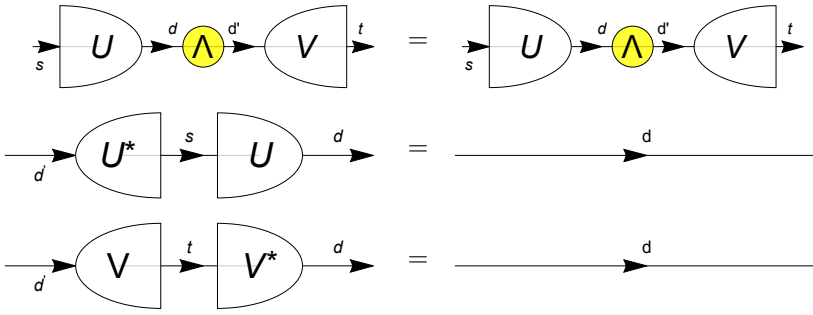

(3.10)

(3.11)

(3.12)

I used these curved shapes as a mnemonic about the orthonormality of the columns or rows of the U or V. The idea is that they are two half-circles, that get put together to make a circle – which is what I am using for a diagonal matrix.

For a Hermitian matrix, the singular value decomposition is identical to the eigenvalue decomposition. Generically, however, they are different, and hold different meanings.

The way I like to think about unitary matrices, is that they are coordinate transforms. They take one orthonormal basis, and transform it into another, preserving all lengths and angles. In this way of thinking, a rectangular row-orthogonal matrix, such as V or $U^\dagger$ is a projector into a lower dimensional subspace. The operator $V^\dagger$ is the pull-back into that bigger space. The fact that $VV^\dagger = I$, tells us that $V^\dagger$ is an embedding of the small space into the big space.

What is the meaning of $VV^\dagger$?

The operator $VV^\dagger$ is the projector into the image of V.

The real power of the SVD comes from using it to approximate matrices. If we throw away the smallest singular values, we develop an approximation to our original matrix. We can see how this works with images. We can interpret this image of a cat as a matrix with $600 \times 400 = 240,000$ entries:

If we do a SVD, and keep only the 20 largest singular values, we need only store $600 \times 20 + 400 \times 20 + 20 = 20,020$ numbers, and get a reasonable cat:

The reason that this works is that the rows and columns of the cat picture are correlated.

D.1. Relationship to other decompositions

One useful relationship is that one can relate the singular value decomposition to the eigenvalue decomposition of $M^\dagger M$ and $MM\dagger$:

$$M^\dagger M = (V^\dagger \Lambda U^\dagger)(U\Lambda V) = V^\dagger \Lambda^2 V \tag{3.13}$$

$$MM^\dagger = (U\Lambda V)(V^\dagger \Lambda U^\dagger) = U\Lambda^2 U^\dagger. \tag{3.14}$$

Thus the singular values are the square roots of the eigenvalues of $M^\dagger M$ and $MM^\dagger$. The rows of V are the eigenvectors of $M^\dagger M$, while the rows of $U^\dagger$ are the eigenvectors of $MM^\dagger$.

The other common connection is to the "Polar Decomposition" – the matrix equivalent of writing $z = re^{i\theta}$. In particular

$$M = U\Lambda V = (UV)(V^\dagger \Lambda V) = (U\Lambda U^\dagger)(UV). \tag{3.15}$$

The combination UV is unitary, and the others are Hermitian (and even already diagonalized).

E. Entanglement

Suppose you divide a quantum system into two parts, A and B. Consider an arbitrary wavefunction $|\psi\rangle_{AB}$. The Schmidt decomposition of ψ is

$$|\psi\rangle_{AB} = \sum_j \lambda_j |\phi^j\rangle_A |\theta^j\rangle_B, \tag{3.16}$$

with

$$\langle\phi^i|\phi^j\rangle = \delta_{ij} \tag{3.17}$$
$$\langle\theta^i|\theta^j\rangle = \delta_{ij}. \tag{3.18}$$

This is nothing but the Singular Value Decomposition. The singular values, λ_j, are also referred to as the Schmidt values. If the original state was properly normalized then

$$\sum_j \lambda_j^2 = 1. \tag{3.19}$$

The Schmidt values tells you about how much subsystem A cares about subsystem B. If there is only one non-zero Schmidt value, you have a product state, and the two parts of the system are independent. Otherwise the parts are entangled. One quantifies the entanglement via the *entanglement entropy*

$$S_{ab} = -\sum_j \lambda_j^2 \log \lambda_j^2. \tag{3.20}$$

Another way to understand this decomposition, is to trace out part of the system:

$$\rho_A = \mathrm{Tr}\Big[|\psi\rangle_{AB}\ {}_{AB}\langle\psi|\Big] = \sum_\nu\ {}_B\langle\nu|\psi\rangle_{AB}\ {}_{AB}\langle\psi|\nu\rangle_B, \tag{3.21}$$

where we have explicitly stated which space each vector lives in. We can do the trace in any basis. In particular, we can use the $|\theta^i\rangle$ as a nice orthonormal basis, to get

$$\rho_A = \sum_j \lambda_j^2 |\phi_j\rangle\langle\phi_j|. \tag{3.22}$$

We therefore see that λ_j^2 is the j'th eigenvalue of the reduced density matrix – and $|\phi_j\rangle$ is the associated eigenvector. This is the same construction we found for relating the SVD of M to the eigenstates of $M^\dagger M$. We can similarly get $|\theta_j$ as the j'th eigenvector of the reduced density matrix for subsystem B.

We therefore see that the entanglement entropy is the entropy of the ensemble generated by tracing over part of the system. We will find that tensor networks are a natural way to encode the geometry of entanglement in a quantum state. Further we will find that by making arguments about the structure of entanglement, will lead us to approximations which help us solve the many-body problem.

F. Geometry of Entanglement

Consider a typical bipartition of a system:

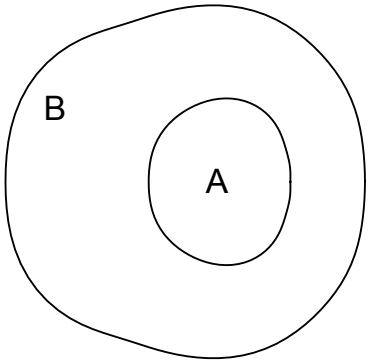

We will imagine that it is a lattice system. We will call V the total number of sites – the "volume" of the system – but it may not be a 3D system. We will call V_A the number of sites in A, and V_B the number of sites in B.

F.1. Generic States

For a generic state – the state of B will depend on all details of A. The Hilbert space of A has a size which is exponential in V_A, so the entanglement entropy, which will be the *log* of this, should scale as

$$S \sim V_A. \tag{3.23}$$

F.2. Gapped Ground States

The situation is very different for a "gapped" ground state. Think for example of a ferromagnet. There is a finite correlation length, and fluctuations in region B, only depend on the parts of A which are within ξ of the boundary. Therefore the entropy should scale as

$$S \sim \Omega_A \xi, \tag{3.24}$$

where Ω is the surface area/ perimeter of the region. This is known as the "Surface law."

Now you might argue that gapped states are boring. While there is some truth to that claim, there is actually a lot of interesting physics with gapped states. In fact, these are conventionally "ordered" state – like magnets and superconductors. We will develop tools which take advantage of the surface-law scaling of gapped states in 1D. Of course, the "surface" in 1D is just a pair of points. So in 1D with a gapped system, the entanglement entropy is independent of the size of the region you are looking at. This is good for doing numerics.

F.3. Free Particles

For non-interacting particles, the only way that the states of subsystem A and B are correlated is by number conservation.

Suppose region A has V_A sites, and the system has an average density of n particles per site. The mean number of particles in region A is $N_A = nV_A$. The fluctuations are $\sqrt{N_A}$. Therefore one expects the entanglement entropy to scale as

$$S_A \sim \log \sqrt{N_A} = \frac{1}{2} \log V_A + \cdots \tag{3.25}$$

Free particles (at fixed density) have an entropy which scales as the logarithm of the subsystem.

On the flip side, if one instead has N total particle, and a total volume V, and divide the system in half. By the same argument

$$S \sim \sqrt{N} \tag{3.26}$$

is independent of system size.

F.4. Topological Order

There is a special case of gapped states, namely those which are "topologically ordered". Topological order was a concept which I believe was first developed in thinking about spin systems, and contemplating ground states which had no local order parameter, but were gapped. The name "topological" has a number of meanings: First, these systems can be distinguished from ordinary "topologically trivial" states, and from one-another, by asking if one can continuously deform one into another without opening a gap. Topology in the mathematics of continuity. There is another way that "topology" enters into topologically ordered states, namely that the ground-state degeneracy depends on the topology of the space. These ground states cannot be distinguished by any local measurement.

A related feature of topologically ordered states is that there exist bulk excitations which cannot be created by any local operator. There may, however, be local operators which can create a finite number of them. The ground state degeneracy is related to the properties of these "quasiparticles." If you create a pair of them, move one around a non-contractable loop, then recombine them, you typically move in the space of ground states.

If you have a topologically ordered state with boundaries, there will generically be gapless edge states. The hand-waving picture is that you are spatially moving between two topologically distinct regions – which is only possible if you locally close a gap.

There are also weaker flavors of "topological order," where you ask about continuously connecting states without breaking any symmetries. These are interesting, but do not necessarily have all of the features of the more generic kind of topological order.

In the context of these lectures, the interesting thing about topologically ordered states is that there is some sort of very diffuse entanglement in the system: The degenerate ground states are distinguished by some "loop-correlations" which extend over non-contractible spatial loops. Consequently, in a bipartition, subsystems B and A are entangled not only by physics happening at the boundary, but also by this non-local

physics. This gives a sub-dominant contribution to the entanglement entropy:

$$S \sim \Omega_A \xi - \gamma, \tag{3.27}$$

where γ is a constant which is independent of the size of A. Of course there is some ambiguity here: How exactly do you define the area? If you use two definitions which differ by an additive constant, then that changes the perceived γ. Further there are corrections due to "corners" and other features of the boundary geometry. There are also some more technical subtleties: For example, if A is non-contractable, then the entanglement entropy actually depends on which ground state you are in. Nonetheless physicists have figured out those details, and one can actually use the scaling of the entanglement entropy as a mechanism for distinguishing topological phases. Our tensor network ideas give a framework for thinking about this entanglement.

Quantum Info

There is an equivalent definition of topological order which is to say that two wavefunctions are in the same topological class if a finite depth quantum circuit can transform one into the other. Topologically distinct wavefunctions cannot be connected by a finite depth circuit. Symmetry protected phases add constraints on the gates.

F.5. Critical States

Critical systems have correlation functions which fall off as power laws. In 1D this leads to a

$$S_A \sim \log V_A. \tag{3.28}$$

In higher dimensions there is typically an area law with logarithmic contribution. So in 2D, you would get $S \sim L \log L$, where L is the perimeter.

Complete the fourth intro to tensors notebook – which concludes that introduction.

G. Homework 3 – Due Feb 5

Problem 12. For Credit: Take any image. Use the SVD on it. Make a plot of the singular values λ_j as a function of j. Truncate to a small number of singular values. How many do you need to keep for the image to look good? By how much did you compress the image. [Note that you won't do as well as *.jpg*, unless you do an image which has a lot of horizontal and vertical lines – like a flag or a Mondrian painting.]

Problem 13. For Credit A 'cat' state is $|\psi\rangle = |\uparrow\uparrow\uparrow\uparrow\uparrow\rangle + |\downarrow\downarrow\downarrow\downarrow\downarrow\rangle$. Find its entanglement entropy if it is split in two

Problem 14. For Credit The "W"-state is $|\uparrow\downarrow\downarrow\downarrow\downarrow\rangle + |\downarrow\uparrow\downarrow\downarrow\downarrow\rangle + |\downarrow\downarrow\uparrow\downarrow\downarrow\rangle + |\downarrow\downarrow\downarrow\uparrow\downarrow\rangle + |\downarrow\downarrow\downarrow\downarrow\uparrow\rangle$. Find its entanglement entropy if it is split between the second and third state.

Problem 15. Bonus

15.1. Write a program which will take in the matrix describing a uniform infinite matrix product state, and give you the transfer matrix, and its eigenvalues. Test it with with the wavefunctions from section A of this chapter. Verify the single-particle wavefunction yields a non-diagonalizable transfer matrix, and the chiral state yields one with complex eigenvalues.

15.2. What physical property does the single particle state have that makes the transfer matrix non-diagonalizable?

15.3. What physical property does the chiral state have that gives the transfer matrix complex eigenvalues?