Quantum eigenvalue transformation and spectral filtering

Use spectral filters to retain useful eigenvalue regions while tracking gaps and success probability.

Research 21 min lesson Dr. Charith Jayasekara

By the end, you'll be able to

- Compute the success of a two-level spectral projection.
- Distinguish filter accuracy from initial overlap.
- Explain why a sharp threshold needs a spectral promise.

Select directions by their eigenvalues

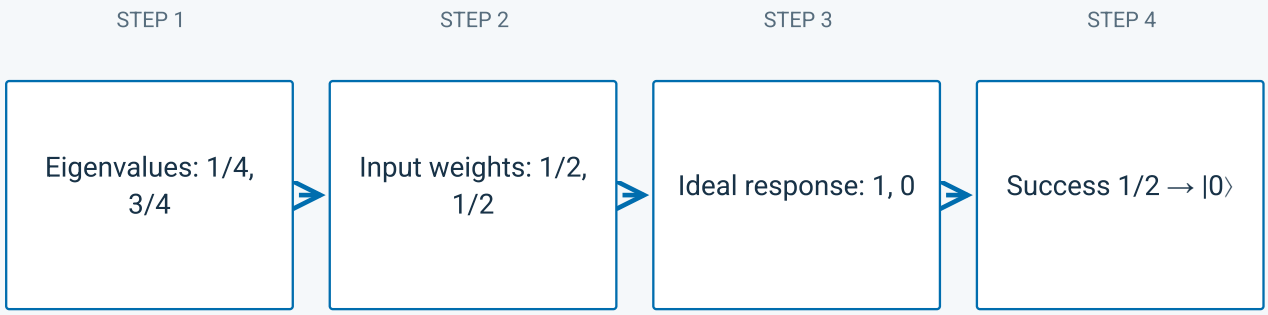
An eigenvector is a direction that a matrix scales without mixing with other directions. Spectral filtering changes the weights of those directions according to their eigenvalues. If a Hamiltonian's low-energy directions represent useful states, a filter can suppress unwanted energies. If a data matrix has a few important spectral components, a filter can retain those components while damping others.

The algorithmic challenge is to apply that response without first printing the entire matrix spectrum. Quantum methods can work coherently on a superposition of eigenvectors. The output is usually a transformed state or a measurement probability, not a classical list of all eigenvectors. That output task description is essential when interpreting a claimed improvement.

Review linear algebra, normalisation and measurement. Advanced readers can connect this chapter to polynomial approximation and the spectral theorem for Hermitian matrices.

WORKED DIAGRAM · AA05

Keep the lower-energy direction



An equal input filtered onto the 1/4 eigenspace has unnormalised amplitude 1/√2 and success probability 1/2. Normalising the selected state does not make selection free.

State the matrix-access setup

Let H be Hermitian, with known spectral range after rescaling. Suppose we can implement a block encoding or another efficient coherent access primitive for H. Write the input as a sum of eigenvectors with coefficients c_j. The desired operation weights each component by a function f of its eigenvalue. The resulting unnormalised state is a sum of c_j times f(lambda_j) times the eigenvector.

The function must be implemented within a unitary process, often using an ancilla and conditioning on a selected outcome. A contraction bounded in magnitude by one fits naturally as a block. The success probability depends on the norm of the filtered state. Renormalising a tiny surviving component is not free: it can require repeated preparation or amplitude amplification.

Work a two-level filter

Take H diagonal with eigenvalues one quarter and three quarters. Prepare an equal superposition of its two basis eigenvectors. Suppose an ideal conceptual filter has response one on the lower eigenvalue and zero on the higher. The selected output is the zero basis state divided by square root of two. Its squared norm is one half, so the conditioning succeeds with probability one half. Conditional on success, the state is zero.

Now suppose the input had only one percent probability weight on the low-energy direction. The same ideal filter would succeed only one percent of the time. A theorem about implementing the filter does not guarantee that the initial state overlaps usefully with the target subspace. Preparing a better initial state or amplifying success must be included in the resource analysis.

An actual finite-degree polynomial filter approximates the ideal step over promised separated intervals. It may leave small unwanted components and slightly reduce the desired component. Quantify those errors before normalisation. If the desired overlap is very small, a small absolute filtering error can become a significant error in the conditional state.

Write both the selected branch and its probability before normalisation.

$$H = \mathrm{diag}(1/4, 3/4), \quad f(1/4) = 1,\ f(3/4) = 0, \qquad f(H)|+\rangle = \frac{|0\rangle}{\sqrt{2}}, \quad P_{\text{success}} = \frac{1}{2}.$$

This is an ideal filter defined on the two eigenvalues. Conditioning yields $|0\rangle$. An actual polynomial approximates the required separated spectral intervals, and low initial overlap can make success rare.

Two ways to reveal spectral information

Phase estimation can attach approximate eigenvalue information to an auxiliary register. A conditional rotation can then weight components according to that information, followed by uncomputing the estimate. Polynomial methods instead apply a designed spectral response through a block encoding without necessarily storing a readable eigenvalue register. Both approaches require controlled access, precision and a success analysis.

The QSVT framework provides systematic polynomial transformations of singular values and, under suitable Hermitian encodings and polynomial conventions, eigenvalue-related transformations. Its bounded-polynomial constraints and parity rules are part of the method. It does not permit arbitrary discontinuous maps at no cost. Read the QSVT framework.

Follow a filtering workflow

First decide what spectral region counts as useful. State a gap or transition region so the filter does not need to resolve arbitrarily close eigenvalues. Rescale the matrix and construct the access operation. Choose a bounded approximate response and an error budget. Apply the corresponding coherent sequence, measure or amplify the selection flag, and evaluate the observable required by the application.

Check the initial overlap and output normalisation separately from polynomial accuracy. A filter that approximates its target uniformly can still be a poor state-preparation method if nearly all initial weight is discarded. Likewise, an efficiently prepared filtered state may not support cheap extraction of every coordinate. Match the final measurement to the scientific or computational question.

Why a gap matters

Suppose two eigenvalues lie on opposite sides of a threshold but are extremely close. Distinguishing them requires high spectral resolution. A sharp filter must therefore use a more demanding polynomial or longer coherent evolution. A promised separation makes an efficient approximate filter possible under favourable conditions. The gap is a resource parameter, not a minor technical footnote.

Filtering also differs from measuring energy and keeping a classical estimate. A coherent filter can preserve superpositions within the retained subspace, whereas an overly informative measurement may distinguish its components. Decide whether the application needs a state, a sampled energy label or an expectation value. These are related but different tasks.

Practice with answers

Exercise 1. If an exact filter retains an input component with probability weight 0.09, what is the selected output norm and success probability?

Check the worked answer

The unnormalised output norm is 0.3, and its squared norm, 0.09, is the selection probability. Conditional normalisation gives unit norm but does not increase how often the condition succeeds.

Exercise 2. Why is “prepare the ground state by applying a low-energy filter” not a complete efficient algorithm?

Check the worked answer

It needs an implementable access operation, a spectral separation and accuracy analysis, adequate initial overlap, a selection or amplification cost, and a useful output measurement. A filter description specifies a desired transformation; it does not establish that all those resources are small.

Check your understanding

PAUSE & THINK

If a filter retains a component of probability weight 0.09, what is its success probability?

- ☐ 0.3
- ☐ 0.09
- ☐ 1

Answer: 0.09. The selected output norm is 0.3, and its square 0.09 is the success probability. Conditional normalisation does not alter how often success occurs.

Sources

- Read the QSVT framework