Structured Decomposition of the Quantum Atypical Sector

A. Chawla

April 2, 2026

Abstract

This note develops the second step of the structured quantum source-coding program. Starting from the typical–atypical decomposition established in Step 1, it introduces spectral probability zones, defines a structured orthogonal decomposition of the atypical subspace, isolates the dominant atypical component, and records the asymptotic conditions under which the additional structure remains consistent with Schumacher compression. The purpose of the present step is not yet to choose a clustering geometry, but rather to identify the formal objects and accounting relations that any later structured encoder must respect.

1 Introduction

In weak source coding, the principal modification of Shannon's classical scheme is to replace the single atypical error declaration by a partition of the atypical set into classes, zones, or residual components. The purpose of the present step is to formulate the quantum analogue of that move.

Step 1 established the spectral representation of a memoryless quantum source, the typical subspace $T_\epsilon^{(n)}$, the atypical projector

$$Q_\epsilon^{(n)} = I - \Pi_\epsilon^{(n)},$$

and the asymptotic Schumacher compression rate $S(\rho)$.

It also identified the atypical sector as a low-weight spectral tail with finite-blocklength internal structure. From the standpoint of the larger program, this observation is decisive: once the atypical sector is seen not as a featureless remainder but as a structured tail in the spectrum of $\rho^{\otimes n}$, it becomes natural to ask how much of that structure can be retained and exploited.

The role of the present step is therefore organizational. It does not yet claim a gain in compression fidelity, nor does it prescribe a unique clustering rule. Instead, it introduces the formal language in which such refinements can later be stated. In particular, it identifies the relevant probability zones, the projectors attached to those zones, the possible decompositions of the atypical sector, and the weights of the resulting components.

The role of this step is threefold:

1. To introduce the three-zone quantum counterpart of the weak-coding picture.
2. To define orthogonal atypical components together with their weights and dominant contribution.
3. To record the structural overhead needed to distinguish these components while preserving asymptotic consistency with Schumacher compression.

No geometric criterion is imposed at this stage.

This deferral is intentional. Step 3 will determine how the decomposition should be chosen by introducing distance-based clustering in the quantum state space. Step 4 will explain how the resulting component structure enters a rate–fidelity tradeoff. The present note should therefore be read as the formal bridge between the spectral typicality of Step 1 and the operational refinements that follow.

2 Quantum Probability Zones

Let

$$\rho = \sum_{i=1}^{d} \lambda_i \left|e_i\right\rangle \left\langle e_i\right|$$

be the spectral decomposition of the source state, and let

$$\rho^{\otimes n} = \sum_{i_1,\dots,i_n} \lambda_{i_1} \cdots \lambda_{i_n} \left|e_{i_1} \cdots e_{i_n}\right\rangle \left\langle e_{i_1} \cdots e_{i_n}\right|$$

be the induced block source.

The key observation is that Step 1 already supplies a natural scalar quantity with which the spectrum may be stratified, namely the information density. Once the eigenvalue products of $\rho^{\otimes n}$ are written in logarithmic form, one obtains an immediate analogue of the classical picture in which typicality is defined by comparison with the entropy.

For a spectral word $\mathbf{i} = (i_1, \dots, i_n)$, define its information density by

$$\imath_n(\mathbf{i}) := -\frac{1}{n} \log_2(\lambda_{i_1} \cdots \lambda_{i_n}).$$

The entropy level $S(\rho)$ therefore plays the same organizing role here that $H(X)$ plays in the classical asymptotic equipartition property. The natural next step is to separate the spectral words according to whether their information density is near, above, or below this benchmark.

Definition 1 (Quantum Probability Zones)**.** Fix $\epsilon > 0$.

a) The *Quantum Medium Probability Zone* (QMPZ) is the set of spectral words satisfying

$$|\imath_n(\mathbf{i}) - S(\rho)| \le \epsilon.$$

b) The *Quantum Very Low Probability Zone* (QVLPZ) is the set of spectral words satisfying

$$\imath_n(\mathbf{i}) > S(\rho) + \epsilon.$$

c) The *Quantum Very High Probability Zone* (QVHPZ) is the set of spectral words satisfying

$$\imath_n(\mathbf{i}) < S(\rho) - \epsilon.$$

The same partition may be written directly in terms of eigenvalue bands:

$$\begin{aligned}
\mathbf{i} \in \text{QMPZ} &\iff 2^{-n(S(\rho)+\epsilon)} \le \lambda_{i_1} \cdots \lambda_{i_n} \le 2^{-n(S(\rho)-\epsilon)}, \\
\mathbf{i} \in \text{QVLPZ} &\iff \lambda_{i_1} \cdots \lambda_{i_n} < 2^{-n(S(\rho)+\epsilon)}, \\
\mathbf{i} \in \text{QVHPZ} &\iff \lambda_{i_1} \cdots \lambda_{i_n} > 2^{-n(S(\rho)-\epsilon)}.
\end{aligned}$$

This three-zone language is useful for two reasons. First, it makes the connection with weak source coding explicit. Second, it separates two conceptually different kinds of atypicality: spectral words may be atypical because they are too rare relative to the entropy scale, or because they are too concentrated. Later geometric clustering can then be asked to operate within these tails rather than on an undifferentiated complement.

Remark 1. QMPZ is exactly the spectral typical region of Step 1.

The atypical sector decomposes into the two tails QVLPZ and QVHPZ. This is the direct quantum counterpart of the medium-, very-low-, and very-high-probability zones used in weak source coding.

3 Projectors Onto the Three Zones

Once the zones have been defined at the level of spectral words, the corresponding subspaces and projectors are immediate. This step is formally simple, but it is important because it converts the three-zone classification into Hilbert-space objects that can enter later encoder and decoder constructions.

Definition 2 (Zone Projectors)**.** Define the orthogonal projectors

$$\begin{aligned}\Pi_{\mathrm{M}}^{(n)} &:= \sum_{\mathbf{i}\in\mathrm{QMPZ}} |e_{\mathbf{i}}\rangle\langle e_{\mathbf{i}}|,\\ \Pi_{\mathrm{L}}^{(n)} &:= \sum_{\mathbf{i}\in\mathrm{QVLPZ}} |e_{\mathbf{i}}\rangle\langle e_{\mathbf{i}}|,\\ \Pi_{\mathrm{H}}^{(n)} &:= \sum_{\mathbf{i}\in\mathrm{QVHPZ}} |e_{\mathbf{i}}\rangle\langle e_{\mathbf{i}}|,\end{aligned}$$

where $|e_{\mathbf{i}}\rangle = |e_{i_1}\cdots e_{i_n}\rangle$.

These projectors resolve the identity:

$$\Pi_{\mathrm{M}}^{(n)} + \Pi_{\mathrm{L}}^{(n)} + \Pi_{\mathrm{H}}^{(n)} = I, \qquad \Pi_{\mathrm{M}}^{(n)}\Pi_{\mathrm{L}}^{(n)} = \Pi_{\mathrm{M}}^{(n)}\Pi_{\mathrm{H}}^{(n)} = \Pi_{\mathrm{L}}^{(n)}\Pi_{\mathrm{H}}^{(n)} = 0.$$

Moreover,

$$\Pi_{\mathrm{M}}^{(n)} = \Pi_{\epsilon}^{(n)}, \qquad Q_{\epsilon}^{(n)} = \Pi_{\mathrm{L}}^{(n)} + \Pi_{\mathrm{H}}^{(n)}.$$

Thus the typical projector of Step 1 is recovered as the projector onto QMPZ, while the atypical projector is now split into two spectrally distinguished pieces. This makes precise the statement that the atypical sector already contains a nontrivial internal organization before any clustering rule has been imposed.

Proposition 1 (Weight of the Quantum Zones)**.** *Let*

$$w_{\mathrm{M}}^{(n)} = \mathrm{Tr}\Big(\rho^{\otimes n}\Pi_{\mathrm{M}}^{(n)}\Big), \quad w_{\mathrm{L}}^{(n)} = \mathrm{Tr}\Big(\rho^{\otimes n}\Pi_{\mathrm{L}}^{(n)}\Big), \quad w_{\mathrm{H}}^{(n)} = \mathrm{Tr}\Big(\rho^{\otimes n}\Pi_{\mathrm{H}}^{(n)}\Big).$$

Then

$$w_{\mathrm{M}}^{(n)} + w_{\mathrm{L}}^{(n)} + w_{\mathrm{H}}^{(n)} = 1,$$

and, for sufficiently large n,

$$w_{\mathrm{L}}^{(n)} + w_{\mathrm{H}}^{(n)} \le \epsilon.$$

Proof. The first identity follows from the orthogonal resolution of the identity.

For the second, Step 1 gives

$$w_{\mathrm{L}}^{(n)} + w_{\mathrm{H}}^{(n)} = \mathrm{Tr}\Big(\rho^{\otimes n}Q_{\epsilon}^{(n)}\Big) \le \epsilon$$

for sufficiently large n. □

The proposition shows that the medium zone carries asymptotically almost all of the spectral weight, while the union of the two tails remains small. Nevertheless, the tails need not be negligible for finite blocklength. It is precisely in this finite-n regime that a structured treatment of the atypical part may become meaningful.

4 Structured Decomposition of the Atypical Subspace

The quantum atypical subspace is

$$B_\epsilon^{(n)} := \mathrm{Ran}(Q_\epsilon^{(n)}).$$

Up to this point, the atypical sector has been split only into the two coarse tails QVLPZ and QVHPZ. The next definition introduces the more general object that will be used in later steps: an arbitrary orthogonal resolution of the atypical subspace into components.

Definition 3 (Structured Atypical Decomposition)**.** A *structured atypical decomposition* of $B_\epsilon^{(n)}$ is a family of mutually orthogonal subspaces

$$C_1^{(n)}, \dots, C_k^{(n)} \subseteq B_\epsilon^{(n)}$$

such that

$$B_\epsilon^{(n)} = \bigoplus_{r=1}^{k} C_r^{(n)}.$$

If $\Pi_r^{(n)}$ denotes the projector onto $C_r^{(n)}$, then

$$Q_\epsilon^{(n)} = \sum_{r=1}^{k} \Pi_r^{(n)}.$$

This definition is intentionally broad. Its purpose is not to decide which decomposition is optimal, but to identify the class of decompositions that a structured theory may consider. At this level one wants a formal framework with enough flexibility to accommodate spectral shells, zone refinements, or later metric-induced clusters.

Remark 2. The decomposition is intentionally left abstract.

It may simply separate QVLPZ from QVHPZ, it may refine one of the two tails, or it may partition the tails into shells or smaller components. In Step 3, the choice of decomposition will be guided by metric-based clustering in the quantum state space.

Proposition 2 (Component Weight Identity)**.** *For a structured atypical decomposition, define*

$$w_r^{(n)} := \mathrm{Tr}\left(\rho^{\otimes n} \Pi_r^{(n)}\right), \qquad r = 1, \dots, k.$$

Then

$$\sum_{r=1}^{k} w_r^{(n)} = \mathrm{Tr}\left(\rho^{\otimes n} Q_\epsilon^{(n)}\right) \le \epsilon$$

for sufficiently large n.

Proof. Since $Q_\epsilon^{(n)} = \sum_{r=1}^{k} \Pi_r^{(n)}$ and the trace is linear,

$$\sum_{r=1}^{k} w_r^{(n)} = \sum_{r=1}^{k} \mathrm{Tr}\left(\rho^{\otimes n} \Pi_r^{(n)}\right) = \mathrm{Tr}\left(\rho^{\otimes n} Q_\epsilon^{(n)}\right).$$

The upper bound follows from Proposition 1. $\square$

This identity is the basic accounting relation of the structured theory. It ensures that every atypical spectral component is counted exactly once and that no hidden remainder survives outside the chosen decomposition. Later encoder constructions can therefore allocate resources component by component without ambiguity.

5 Dominant Atypical Component

Once a decomposition has been chosen, the next natural question is whether one of its components carries most of the atypical mass. The answer to this question will later determine whether a differentiated treatment of the atypical sector is worthwhile.

Definition 4 (Dominant Atypical Component)**.** For a structured atypical decomposition, define the dominant atypical weight by

$$D_Q^{(n)} := \max_{1 \le r \le k} w_r^{(n)}.$$

Any component attaining this maximum will be called a *dominant atypical component.*

The quantity $D_Q^{(n)}$ is the quantum analogue of the largest residual class weight in weak source coding. It is important because any refined encoder that singles out a privileged atypical component will inevitably be controlled by this maximal weight rather than by the full atypical weight alone.

Proposition 3 (Elementary Bounds on the Dominant Weight)**.** *For every structured atypical decomposition,*

$$\frac{1}{k} \operatorname{Tr}\left(\rho^{\otimes n} Q_\epsilon^{(n)}\right) \le D_Q^{(n)} \le \operatorname{Tr}\left(\rho^{\otimes n} Q_\epsilon^{(n)}\right).$$

Consequently, for sufficiently large n,

$$D_Q^{(n)} \le \epsilon.$$

Proof. Since $D_Q^{(n)}$ is the maximum of the nonnegative numbers $w_1^{(n)}, \ldots, w_k^{(n)}$,

$$k D_Q^{(n)} \ge \sum_{r=1}^{k} w_r^{(n)}$$

and also

$$D_Q^{(n)} \le \sum_{r=1}^{k} w_r^{(n)}.$$

Using Proposition 2 gives the stated inequality. The final statement follows from

$$\operatorname{Tr}\left(\rho^{\otimes n} Q_\epsilon^{(n)}\right) \le \epsilon.$$

□

The proposition does not yet prove that a structured encoder helps. What it does provide is the correct quantitative target. If later geometric clustering succeeds in concentrating most of the atypical mass into a comparatively small number of components, then $D_Q^{(n)}$ will be the quantity through which that concentration becomes operationally visible.

6 Structural Overhead Parameter

Weak source coding tracks the number of atypical classes through an explicit structural parameter. The corresponding quantum parameter is as follows.

Definition 5 (Quantum Structural Overhead)**.** If the atypical sector is decomposed into k orthogonal components, define

$$\chi_Q := \frac{\log_2 k}{n(S(\rho) + \epsilon)}.$$

The meaning of χ_Q is straightforward: it measures the relative overhead required to specify which atypical component has occurred, using the Schumacher baseline scale $n(S(\rho)+\epsilon)$ as the reference denominator. Thus χ_Q does not yet represent a theorem about rate improvement or deterioration. Rather, it is the natural bookkeeping parameter against which later operational gains must be compared.

Remark 3. The quantity χ_Q measures the relative rate overhead required to label k atypical components with respect to the Schumacher baseline scale $n(S(\rho)+\epsilon)$.

If a later encoder attaches a component label to atypical states, then $\log_2 k$ is the number of additional label bits, and χ_Q is the normalized form of that overhead.

Remark 4. The number of components k may be fixed in advance or may grow with the blocklength.

A coarse decomposition corresponds to small k, whereas a fine decomposition may allow k to increase with n. If

$$\log_2 k = o\big(n(S(\rho)+\epsilon)\big),$$

then

$$\chi_Q \to 0 \qquad \text{as} \qquad n \to \infty,$$

so the relative overhead vanishes asymptotically.

7 Asymptotic Consistency

The structured decomposition introduced in this step must remain compatible with Schumacher compression in the asymptotic limit. This requirement is conceptually important: the structured theory is intended as a refinement of Schumacher coding, not as a different asymptotic coding law.

Proposition 4 (Asymptotic Recovery of Schumacher Compression)**.** *Suppose $\epsilon = \epsilon_n \to 0$ and the number of atypical components satisfies*

$$\log_2 k_n = o\big(n(S(\rho)+\epsilon_n)\big).$$

Then

$$\chi_Q^{(n)} \to 0 \qquad \textit{and} \qquad D_Q^{(n)} \to 0$$

as $n \to \infty$.

In this sense, the structured refinement reduces asymptotically to Schumacher compression.

Proof. The first claim is immediate from the definition

$$\chi_Q^{(n)} = \frac{\log_2 k_n}{n(S(\rho)+\epsilon_n)}.$$

For the second, Proposition 3 yields

$$0 \le D_Q^{(n)} \le \operatorname{Tr}\Big(\rho^{\otimes n} Q_{\epsilon_n}^{(n)}\Big) \le \epsilon_n$$

for sufficiently large n, hence $D_Q^{(n)} \to 0$. □

Remark 5. The preceding proposition records explicitly what was only implicit in Step 1.

The structural overhead and the dominant atypical weight can both be made asymptotically negligible, provided the component count grows slowly enough. This ensures that Step 2 remains a refinement of Schumacher coding rather than a competing asymptotic theory.

8 Example: Diagonal Qubit Source

It is useful to see the preceding constructions in a concrete source for which the spectral structure is completely explicit. The diagonal qubit source provides exactly such an example.

Let

$$\rho = \mathrm{diag}(1-p, p), \qquad 0 < p < 1.$$

For a spectral word $\mathbf{i} \in \{1,2\}^n$, let $m(\mathbf{i})$ denote the number of occurrences of the symbol 2. Then

$$\lambda_{\mathbf{i}} = (1-p)^{n-m(\mathbf{i})} p^{m(\mathbf{i})}.$$

All words with the same value of m have the same eigenvalue, so the block source admits the shell decomposition

$$\rho^{\otimes n} = \sum_{m=0}^{n} (1-p)^{n-m} p^m P_m^{(n)},$$

where

$$P_m^{(n)} := \sum_{\substack{\mathbf{i} \in \{1,2\}^n \\ m(\mathbf{i})=m}} |e_{\mathbf{i}}\rangle \langle e_{\mathbf{i}}| .$$

The corresponding shell weight is

$$\mathrm{Tr}\Big(\rho^{\otimes n} P_m^{(n)}\Big) = \binom{n}{m} (1-p)^{n-m} p^m.$$

Hence the source is already stratified by empirical frequency before any clustering algorithm is introduced. The shells provide a natural finite-blocklength language for discussing typicality, atypicality, and later component grouping.

Proposition 5 (Zone Description for the Diagonal Qubit Source)**.** *For the diagonal qubit source, the three quantum probability zones are determined by the empirical frequency m/n:*

$$\begin{aligned}
P_m^{(n)} \subseteq \textsc{qmpz} &\iff \left| -\frac{1}{n} \log_2\!\big((1-p)^{n-m} p^m\big) - h_2(p) \right| \le \epsilon, \\
P_m^{(n)} \subseteq \textsc{qvlpz} &\iff -\frac{1}{n} \log_2\!\big((1-p)^{n-m} p^m\big) > h_2(p) + \epsilon, \\
P_m^{(n)} \subseteq \textsc{qvhpz} &\iff -\frac{1}{n} \log_2\!\big((1-p)^{n-m} p^m\big) < h_2(p) - \epsilon.
\end{aligned}$$

Proof. All basis vectors in the shell indexed by m have the same eigenvalue $(1-p)^{n-m}p^m$.

Hence the information density is constant on that shell and equal to

$$-\frac{1}{n} \log_2\!\big((1-p)^{n-m} p^m\big).$$

The result follows by substituting this value into the defining inequalities for QMPZ, QVLPZ, and QVHPZ. □

Remark 6. For this source, the natural shell decomposition already suggests several choices of k.

One may keep the two tails separate, refine each tail into shell groups, or isolate every atypical shell individually. In the latter case $k = O(n)$, so

$$\chi_Q = O\!\left(\frac{\log n}{n}\right),$$

which still vanishes asymptotically. Step 3 will determine which such decompositions are geometrically meaningful.

9 Forward Role of Step 2

The present step defines:

1. The spectral zones QMPZ, QVLPZ, and QVHPZ.
2. The atypical components $C_r^{(n)}$ and their projectors $\Pi_r^{(n)}$.
3. The component weights $w_r^{(n)}$, the dominant atypical weight $D_Q^{(n)}$, and the structural overhead χ_Q.

These are the objects on which the next stages of the program operate. Step 3 will endow the decomposition with geometric meaning by clustering components whose associated states are close in a suitable metric. Step 4 will then ask how those geometrically chosen components should be encoded.

The improvements in fidelity from refining the dominant atypical component through structured encoding must exceed the label overhead χ_Q for the overall rate to remain competitive with Schumacher compression. Step 4 will formalize this tradeoff.

10 Conclusion

Step 2 formalizes the quantum counterpart of weak atypical-set classification.

The medium-probability region is identified with the Schumacher typical subspace, while the atypical sector is split into very-low- and very-high-probability tails. A structured decomposition of the atypical subspace is then introduced, together with the dominant atypical weight, the structural overhead parameter, and the asymptotic conditions under which the refinement remains consistent with Schumacher compression.

The main contribution of this step is therefore conceptual and structural. It shows that the atypical sector of a quantum source may be treated as an organized collection of weighted components rather than as a single undifferentiated failure event. This is the formal setting in which the later geometric and operational refinements become meaningful.