Issue 1228: The SVD Step in Orthonormalization

2026-05-05

1 At a glance

- **Difficulty: medium.** The paper's SVD step is finite-dimensional linear algebra, and the project already has the main positive-Gram route.
- **Estimated weight: 3–6 theorem interfaces** and **300–700 lines** for the local route; a reusable rectangular complex SVD/polar theorem would likely be **800–1500 lines**.
- **Mathlib/project split:** about **65% Mathlib** and **35% project-local**. Mathlib supplies spectral theory, adjoints, orthonormal bases, and square roots; the project supplies the rectangular $X, \widehat{X}, P$ packaging and estimates.
- **Key Mathlib inputs:** finite-dimensional inner-product spaces, matrix adjoints, Hermitian eigenvectors/eigenvalues, positive semidefinite matrices, and continuous-functional-calculus square roots.[1]

2 Key theorem forms to develop

1. **Rectangular polar theorem.** If X is $m \times d$, $m \leq d$, and $Q = X^\dagger X$, construct $\widehat{X}$ with
$$\widehat{X}\widehat{X}^\dagger = I_m, \qquad X^\dagger \widehat{X} = \sqrt{Q}.$$
2. **SVD adapter.** If $X = U\Sigma V^\dagger$ and $\widehat{X} = UI_{m\times d}V^\dagger$, derive the two identities above from the unitary laws and positivity of the SVD middle factor.

[1] Representative APIs include `Orthonormal`, `Matrix.IsHermitian.eigenvectorBasis`, `Matrix.IsHermitian.eigenvalues`, `CFC.sqrt`, and `CFC.sqrt_unique`.

3. **Projective repair estimate.** For $P_a = \widehat{X}^\dagger T_a \widehat{X}$, prove $P_a P_b = \mathbf{1}_{a=b} P_a$ and the P-versus-Q estimate from the squared-difference bound.

3 The source step

The SVD appears in the orthonormalization proof of [JNV$^+$20].[2] After rounding and rank reduction, the paper has projectors Q_a with $Q = \sum_a Q_a$ and total rank at most d. It defines

$$X_a = \sum_i |a,i\rangle \langle v_{a,i}|, \qquad X = \sum_a X_a,$$

(`[blueprint] def:matrix-decomposition-Q`)

with $X_a = T_a X$ and $Q_a = X^\dagger T_a X$. Since X is $m \times d$ with $m \le d$, the paper takes

$$X = U \Sigma_{m\times d} V^\dagger. \qquad \text{(\texttt{[blueprint] def:svd-of-X})}$$

4 What the SVD supplies

The paper replaces the singular values by 1:

$$\widehat{X} = U I_{m\times d} V^\dagger, \qquad P_a = \widehat{X}^\dagger T_a \widehat{X}. \quad \text{(\texttt{[blueprint] def:projective-P})}$$

It uses three consequences: **coisometry** $\widehat{X}\widehat{X}^\dagger = I$, **square-root coupling** $X^\dagger \widehat{X} = \sqrt{Q}$, and

$$(X - \widehat{X})(X - \widehat{X})^\dagger \le (XX^\dagger - I)^2.$$

The last inequality is the scalar fact $(\sigma - 1)^2 \le (\sigma^2 - 1)^2$ for $\sigma \ge 0$. These facts give

$$Q_a \otimes I \approx_{30\zeta^{1/4}} P_a \otimes I, \qquad \text{(\texttt{lem:P-Q-approx})}$$

and hence $A_a \otimes I \approx_{84\zeta^{1/4}} P_a \otimes I$.[3]

[2] The SVD passage is in `references/ldt-paper/orthonormalization.tex`, lines 774–845.

[3] The consequences are in `references/ldt-paper/orthonormalization.tex`, lines 984–1064, using the auxiliary identities in lines 862–896. The P-versus-Q estimate and final triangle step are in lines 1089–1192.

5 Blueprint and formalization

The blueprint follows the same sequence: matrix decomposition, SVD, definition of P, then the three identities and the P-versus-Q estimate.[4] Lean separates **packaging** from **linear algebra**: the data package stores the needed $X, \widehat{X}, P$ identities, while the substantive algebra is provided either by explicit SVD adapters or by the positive-Gram construction.[5] The version of Mathlib pinned here lacks exactly the rectangular polar theorem above, so the project-local positive-Gram route supplies it.[6]

6 Conclusion

There is **no SVD discrepancy**. The paper needs $\widehat{X}$ with $\widehat{X}\widehat{X}^\dagger = I$ and $X^\dagger \widehat{X} = \sqrt{Q}$, plus the squared-difference estimate. The formalization preserves this mechanism via explicit SVD adapters or the positive-Gram route.

References

[JNV+20] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. Quantum soundness of the classical low individual degree test, 2020. arXiv:2009.12982.

[4] See `blueprint/src/chapter/ch04_projective.tex`, lines 704–785 and 833–1153.

[5] The formal names include svdOfX, QXPLayerData, rectangularSvd_xHat_coisometry, rectangularSvd_xHat_mixed_raw, exists_xHat_of_positive_gram_spectrum, and pQApprox_ofRankReductionSigmaRangePositiveGram; see `MIPStarRE/LDT/MakingMeasurementsProjective/QXPLayer/Core.lean` and `MIPStarRE/LDT/MakingMeasurementsProjective/QXPLayerIdentities/ProjectorApprox.lean`.

[6] The local coisometry helper Matrix.exists_mul_conjTranspose_eq_one_of_card_le is in `MIPStarRE/Quantum/FiniteHilbert.lean`. Possible upstreaming is tracked in GitHub issue #1250.