

Referee Report on “An exact algorithm for $U(N)$ matrix models in the gauge-invariant singlet sector”

Manuscript arXiv:2607.13725

Summary and overall assessment

This paper develops a representation-theoretic Hamiltonian-truncation method for bosonic Hermitian matrix quantum mechanics in the gauge-singlet sector. For a single matrix, the authors use Schur-polynomial states $|R\rangle$, indexed by Young diagrams, while for $D > 1$ they use restricted Schur states $|R, \mathbf{r}, a, b\rangle$. These bases are orthogonal, diagonalize the free oscillator Hamiltonian, and reduce the number of states from that of the full DN^2 -oscillator Fock space to a count governed by partitions or generalized Littlewood–Richardson multiplicities.

The proposed computational chain is

$$\begin{aligned} \text{antinormal ordering} &\longrightarrow \text{Wick permutations} \longrightarrow \text{double cosets} \\ &\longrightarrow \text{stabilizer cosets} \longrightarrow \text{character sums.} \end{aligned}$$

In the one-matrix problem, the factorial Wick sum is first reduced to the double-coset formula (3.12). Each representative produces a loop factor and a permutation τ , after which the projected stabilizer $G_{\tau,n}^1$ reduces a second permutation sum to (3.25). Ordinary symmetric-group character theory then gives the explicit formulas (3.30) and (3.31), and the result is assembled in the reusable form (3.32). The authors implement this construction for the quartic oscillator (2.6), obtaining unnormalized matrix elements that are polynomials in N . Figure 9 compares five low-lying levels at $N = 12, 15, 20, 30$ with the exact fermion reformulation of the one-matrix singlet problem.

For several matrices, the center of $\mathbb{C}[S_n]$ is replaced by the centralizer of the Young subgroup $S_n = S_{n_1} \times \cdots \times S_{n_D}$. Equations (3.39)–(3.55) give the formal analogue of the one-matrix reduction in terms of restricted conjugacy classes and restricted characters. The final expression (3.55) is the result of greatest prospective physical interest, since it would apply to genuinely noncommuting matrix degrees of freedom. The authors candidly identify the computation of restricted characters as the present bottleneck, and no multi-matrix Hamiltonian is actually assembled.

The double-coset/stabilizer-coset reduction is a promising idea. In particular, the separation of group-theoretic data from N and from the physical couplings could be useful when many ranks or couplings are to be scanned. I also find the organization in terms of the exact finite- N Schur basis natural and potentially valuable for nonplanar problems. The displayed one-matrix spectral agreement would ordinarily be encouraging, but it cannot presently validate the published derivation: the antinormal decomposition in Table 2 fails a direct vacuum-expectation-value check for every $N > 1$.

The current manuscript, however, does not yet substantiate its broadest claims. The displayed quartic operator identity is incorrect, so the one-matrix matrix elements and spectrum must first be regenerated from a corrected input. Moreover, the only implementation concerns a model already solved efficiently by the fermion mapping; the multi-matrix formula is neither implemented nor checked in a nontrivial multiplicity sector. The numerical eigenproblem in the non-normalized basis is not specified, equation (3.37) contradicts the inner product (2.32), and the quartic coupling used for Figure 9 cannot be unambiguously reconstructed from the conventions in the paper. There is also a conceptual conflation between changing the gauge group and restricting the dynamical matrices to be traceless. These issues affect correctness, reproducibility, and the claimed physical reach.

Recommendation: major revision. The algorithmic core could merit publication only after Table 2 is corrected, every affected one-matrix result is recomputed, the eigenproblem and index conventions are fixed, a genuine multi-matrix test is supplied (or the claims are narrowed), and the novelty and performance are positioned quantitatively relative to existing Schur, fermion, and matrix-model methods.

Correctness and logical structure

The review of the one-matrix Schur basis is largely sound. Equations (2.21)–(2.25) correctly derive the norm

$$\alpha_R = \prod_{(i,j) \in R} (N - i + j),$$

and the Wick/projector argument gives the expected orthogonality. The main one-matrix reduction is also logically well organized: excitation conservation gives (3.6), the centrality of P_R and P_S justifies the double-coset reduction, and the final class-sum calculation is consistent with the useful identity

$$\text{tr}(P_S \rho) = \text{Dim}_S \chi_S(\rho).$$

Thus the group-theoretic reduction underlying (3.30)–(3.32) is plausible, subject to the requested low-level checks and a precise account of normalization. Its quartic application is not presently correct, however. With the commutators and permutation-trace convention of the paper, Wick’s theorem gives

$$\langle 0 |, 4m^2 \text{tr}(X^4), | 0 \rangle = 2N^3 + N.$$

The right-hand side specified by Table 2 instead has vacuum expectation

$$(2N^3 + N) + (-4N - 8N^3) + (4 + 2)(N^3 + N) = 3N.$$

Here the second bracket is the contribution of the two $A \otimes A^\dagger$ entries, and the last bracket is that of the two $A^{\otimes 2} \otimes (A^\dagger)^{\otimes 2}$ entries; all other displayed nonconstant types have vanishing vacuum expectation. The two sides agree only at $N = 1$ (for example, at $N = 2$ they are 18 and 6). Therefore Table 2 is not an operator identity for matrix-valued oscillators. Since that table is explicitly fed into (3.30)–(3.32) and the numerical pipeline, the antinormal decomposition and every reported one-matrix matrix element and spectrum must be recomputed and independently checked.

The logical status changes in Section 3.2. Restricted Schur states and their norms are established technology, and equations (3.39)–(3.55) outline a natural formal extension. But the argument contains no worked restricted-character example, no check against a direct Fock-space calculation, and no test in a sector with multiplicity labels. This is particularly important because the first displayed free matrix element already routes those labels incorrectly. Formula (3.55) should therefore be treated as an unvalidated proposal until its indices, normalizations, Hermiticity, and finite- N restrictions are checked explicitly.

The complexity analysis is similarly conditional. The sharper coset count used in Section 3.1 is reported in (B.55) as an empirical observation for $n, n' \leq 30$, while the multi-matrix $O(n_1 n_2)$ count is asserted rather than proved. In addition, the cost of producing coset representatives and restricted characters cannot be inferred from the number of cosets alone. These qualifications do not diminish the exact algebraic identities, but they do limit the claimed asymptotic efficiency.

Novelty and significance

The Schur and restricted-Schur bases, their finite- N row restrictions, their norms, and the permutation-centralizer algebra are prior results. The likely new contribution is the particular reduction of interacting oscillator matrix elements through (3.12), (3.25), (3.30)–(3.32), together with the restricted-character analogue (3.55) and its computational packaging. The paper should state this distinction explicitly. In particular, Section 4 should not say that completeness and orthogonality of the $U(N)$ Schur basis are “established in this work”; they are reviewed and rederived here.

The potential significance lies in the multi-matrix problem, where simultaneous diagonalization and the fermion mapping are unavailable. At present, however, precisely this part is not operational. The one-matrix benchmark is intended to validate code but yields no new spectrum and, given the failure of Table 2, does not presently validate the published pipeline. The paper also does not show that the proposed route is faster or reaches a more useful regime than the exact fermion method for the same problem. A convincing publication case therefore requires either a nontrivial two-matrix demonstration or a more limited framing as a one-matrix algorithm with a formal multi-matrix extension.

The holographic interpretation also requires restraint. Schur labels have a precise giant-graviton interpretation in protected sectors of $\mathcal{N} = 4$ super Yang–Mills theory, but a generic bosonic matrix oscillator with the interactions (2.6) or (2.7) is not thereby a model with that dual interpretation. The present calculation contains neither the fermions nor the supersymmetry of BFSS/BMN, and it does not yet solve the interacting traceless multi-matrix sector. The mathematical method may still be useful, but the physics claims should track what has actually been demonstrated.

Major comments

- 1. Table 2 fails an elementary vacuum check and must be corrected before the numerical results can be assessed.**

The mismatch displayed above is a decisive test of the starting operator identity, independent of truncation or representation-theory conventions. It is particularly concerning that the constant $2N^3 + N$ equals the left-hand-side vacuum expectation, even though antinormally ordered terms containing annihilators to the left of creators also have nonzero vacuum expectation. This suggests that those contractions were not handled consistently when the 13 terms were assembled.

The authors should redo the antinormal ordering analytically or with an independently verified symbolic routine and print the corrected coefficient table. At minimum, the identity should be tested for $N = 1, 2, 3$ on the vacuum and on a set of one- and two-particle Fock states. The resulting Schur-basis entries should then be compared at low cutoff with direct oscillator contractions before the Hamiltonian is diagonalized. All Figure 9 spectra and error statements must be regenerated from the corrected table. If the current code nevertheless produced the displayed agreement, the code and manuscript are using different decompositions or conventions; that discrepancy must be identified explicitly. Until these checks are supplied, neither the one-matrix validation nor the claim of exact matrix elements is established.

- 2. The multi-matrix result (3.55) needs a nontrivial validation, or the scope of the paper must be narrowed.**

The abstract and conclusion describe an exact algorithm for bosonic $U(N)$ matrix models generally, but the public implementation and all numerical evidence concern one matrix. In Section 3.2 the restricted characters are assumed as input, and the authors acknowledge that no efficient general

method to obtain them is available. Consequently, the paper does not yet compute a single matrix element, eigenvalue, or observable for the commutator-squared Hamiltonian (3.34).

At minimum, the authors should work out a $D = 2$ example at small cutoff using available restricted-character data. The test should include a sector in which a generalized Littlewood–Richardson multiplicity exceeds one, because otherwise the most delicate indices in (3.37), (3.41), (3.54), and (3.55) are invisible. The resulting entries should be compared with direct oscillator Wick contractions at small N and checked for Hermiticity, specieswise selection rule (3.39), and the predicted polynomial dependence on N . A small spectrum or time-dependent observable for (3.34) would establish that the formal reduction is operational. If such a demonstration cannot be supplied, the title, abstract, and conclusions should clearly describe the work as a one-matrix implementation with an outlined multi-matrix extension, and statements about a new computational window on nonplanar holographic dynamics should be removed.

3. The manuscript must state the metric of the truncated eigenproblem and qualify the polynomial-in- N claim.

The displayed states are orthogonal, not orthonormal. Equations (2.22) and (2.33) give their nontrivial norms, while Section 3 computes covariant matrix elements such as $\langle R|H|S\rangle$. The matrix that may be passed to an ordinary Hermitian eigensolver is therefore

$$H_{RS}^{\text{orth}}(N) = \frac{\langle R|H|S\rangle}{\sqrt{\alpha_R(N)\alpha_S(N)}},$$

with the analogous restricted-Schur expression. Alternatively, one may solve the generalized eigenproblem $Hc = EGc$ with $G_{RS} = \alpha_R\delta_{RS}$. The paper states neither operation, yet repeatedly says that the assembled matrix can simply be diagonalized.

This is also important for the central selling point. The raw matrix elements in (3.32) are polynomials in N , but the orthonormal-basis entries generally contain square roots of the norm polynomials. If the implementation evaluates the raw polynomials and then normalizes at fixed N , the manuscript should say so explicitly, give the exact formula, and describe how diagrams with more than N rows are removed before division by a vanishing norm. The abstract and conclusion should distinguish reusable polynomial covariant matrix elements from the normalized Hamiltonian actually diagonalized.

4. Equation (3.37) has the wrong multiplicity-index Kronecker deltas, and all restricted indices should be audited.

Equation (2.32) states

$$\langle R, \mathbf{r}, a, b | S, \mathbf{s}, c, d \rangle \propto \delta_{RS} \delta_{\mathbf{r}\mathbf{s}} \delta_{ac} \delta_{bd}.$$

Since H_{free} is a scalar on a fixed excitation sector, its matrix element must carry exactly this Gram-matrix structure. Equation (3.37) instead contains $\delta_{bc}\delta_{ad}$. This is a definite contradiction, not a convention choice.

The error makes an independent check of the later index routing essential. The authors should rederive (3.41)–(3.55) with a fixed matrix-unit convention, including the adjoint relation between $P_{ab}^{R,r}$ and $P_{ba}^{R,r}$. It would be helpful to show explicitly how (3.54) reduces to the direct representation-theory identity for $\text{tr}(P_{cd}^{S,s}\rho)$. The text around (3.43) should also not invoke “centrality of the centralizer basis”: the matrix units P_{ab} are not central when multiplicities exceed one. What is used is their commutation with the Young subgroup.

5. The coupling convention and the Figure 9 benchmark are not reproducible as written.

The quartic Hamiltonian in (2.6) and again in (3.1) contains $(g^2/2) \text{tr}(X^4)$, whereas the fermion Hamiltonian in (E.3)–(E.4) contains $g^2 \text{tr}(X^4)$. The validation text and the Figure 9 caption then impose $g^2/2 = m^2/N$, but every panel is labeled $g = 1.0$, independent of N . These statements can only be reconciled by introducing an additional dimensionless coupling convention, which is not defined in the paper.

The revision should denote the actual coefficient of $\text{tr}(X^4)$ by an unambiguous symbol, state its numerical value for every panel, and use the same symbol in the fermion comparison. A table of the five reference eigenvalues and the truncated eigenvalues at each N would make the validation quantitatively checkable. The authors should also give convergence data for at least one substantially different coupling: a single mean-field scaling does not establish rapid convergence for the claimed class of finite-coupling problems. Exact low-cutoff matrix-element comparisons, including the $N = 1$ anharmonic oscillator, would test the algebra more directly than spectra alone.

6. The discussion of $U(N)$ versus $SU(N)$ conflates gauge invariance with removal of the trace oscillator.

For full $N \times N$ adjoint matrices, the central $U(1)$ acts trivially by conjugation. The adjoint invariants under $U(N)$ and $SU(N)$ are therefore the same. A different problem arises when the dynamical matrices themselves are restricted from $\mathfrak{u}(N)$ to traceless $\mathfrak{su}(N)$ matrices: the decoupled trace oscillator is then removed and relations such as $\text{tr} A^\dagger = 0$ make the unrestricted Schur set overcomplete. These are two distinct operations.

Sections 1.2 and 4 should be rewritten accordingly. In particular, the standard BFSS model is naturally formulated with a $U(N)$ gauge symmetry and a decoupled center-of-mass $U(1)$ sector; one often studies the interacting traceless sector after separating that mode. The manuscript’s statement that BFSS and BMN are instead formulated with $SU(N)$ gauge group is too categorical. Any proposed penalty or Lagrange-multiplier construction should be formulated as a projection eliminating trace-mode excitations, with a demonstration that it preserves the desired spectrum. This correction is also needed before claiming direct applicability to the gravitational sector of BFSS/BMN.

7. The novelty and practical advantage need a quantitative comparison with prior methods.

The paper correctly cites the Schur and restricted-Schur literature, but it does not isolate which algebraic step is new relative to known exact correlators, permutation-centralizer methods, cut-and-join constructions, and prior interacting/dilatation-operator calculations in restricted-Schur bases. The discussion in Section 4 that earlier dynamics were essentially free-field or planar is too broad. A focused comparison should identify the different subgroups, character data, and outputs, and explain what (3.30)–(3.32) and (3.55) add.

The computational case also needs benchmarks. For the one-matrix quartic model, the fermion method is exact and appears much less costly than building the Schur Hamiltonian through $\Lambda = 20$. The authors should report basis size, peak memory, preprocessing time, per- N evaluation time, and diagonalization time, and compare with both the fermion solver and direct low-cutoff Wick assembly. For the intended multi-matrix applications, comparison with the cited tensor-network and bootstrap approaches should use common observables or at least common scaling variables. Without such evidence, “efficient” and “rapid convergence” should be restricted to the precise tests shown.

8. The complexity estimates and the executable definition of the algorithm are incomplete.

Equation (3.33) uses a coset count whose sharp form (B.55) is based on a finite search through $n, n' \leq 30$. The multi-matrix $O(n_1 n_2)$ claim is likewise not proved. These statements should be

either established for the relevant permutations or labeled as empirical assumptions. The asserted $O(n)$ cost per Murnaghan–Nakayama character-table entry also needs justification, since the number of admissible rim-hook removals is not obviously uniformly linear. GAP running time cannot be inferred solely from the number of output representatives.

Moreover, the diagrammatic maps

$$(\zeta, \sigma, n, n') \mapsto (\tau, k), \quad (\omega, \alpha, \tau) \mapsto (\rho, c)$$

are central to the method but are not defined in pseudocode or by an index-level algorithm. A reader cannot independently implement (3.30) from the paper alone. The revision should supply precise pseudocode, input/output conventions, and small worked examples. The code should be archived at a tagged commit or DOI, with exact commands, dependency versions, and the data needed to reproduce Figure 9; a mutable repository branch is not an adequate archival reference for an algorithm paper.

9. Several supporting mathematical statements require correction or sharper hypotheses.

Section 2.2 says that conjugation-invariant polynomials coincide with symmetric polynomials in the matrix entries. The correct statement is that one-matrix conjugation invariants are generated by trace invariants and may be identified with symmetric polynomials in the eigenvalues; arbitrary symmetric functions of the N^2 entries are not conjugation invariant. The completeness argument should be stated in this form.

Appendix B.1 presents the decomposition of the regular module into irreducibles and the characterization of primitive idempotents as though they held for every finite-dimensional associative algebra. They require semisimplicity. They are valid for $\mathbb{C}[G]$ over $\mathbb{C}$ by Maschke’s theorem, which is all that is needed here, so the hypotheses should simply be restricted. Appendix D.1 calls Hermitian matrices with the ordinary commutator the real Lie algebra $\mathfrak{u}(N)$, although the ordinary commutator of two Hermitian matrices is anti-Hermitian; the authors should use anti-Hermitian generators or explicitly adopt the physics bracket $-i[A, B]$. These points may not change the final one-matrix formula, but they matter in a paper whose contribution is representation-theoretic exactness.

10. Claims about systematic convergence and scattering go beyond what has been shown.

The assertion that the finite cutoff is “the only approximation” is algebraically true but does not by itself prove spectral convergence for every gauge-invariant observable allowed in (2.5). For a semibounded confining polynomial Hamiltonian, nested oscillator spaces provide a plausible Rayleigh–Ritz scheme, but momentum-dependent interactions, nonconfining directions, and general multi-trace terms require domain and convergence conditions. The paper should state a class of self-adjoint, semibounded Hamiltonians for which the truncation is expected to converge and distinguish that statement from the purely exact evaluation of each finite matrix element.

Likewise, once a finite Hamiltonian matrix is assembled one can compute finite-time transition amplitudes and correlation functions. Section 4’s stronger claim of access to “real-time scattering processes” requires asymptotic states and an S -matrix construction, neither of which is provided for the confining oscillator models studied here. This language should be replaced by finite-time transition amplitudes unless an actual scattering prescription is added.

Minor comments

1. Equation (2.9) lists only diagonal entries of $A^\dagger$ while the text calls P a polynomial in all N^2 entries; the arguments of P should show the full entry set.

2. The graphical sentence following Figure 3 says that the final diagram has a loop if and only if $R = S$. Orthogonality follows from $P_R P_S = \delta_{RS} P_R$; loop formation itself is not the criterion.
3. Figure 2's source omits the backslash in one `\draw` command, and the corresponding leg should be restored. Figure 7 is reported by the TeX compiler to exceed the page height and should be resized or split.
4. The caption of Figure 8 cites equation (3.30) twice where the two alternatives are (3.30) and (3.31).
5. The notation $G_{\tau,n}^2$ in the symmetric one-matrix formula should be checked: the projection is onto $S_{n'}$, so an n' subscript would be less misleading.
6. Appendix B.6 says that finite- N labels have “less than N rows.” The correct condition used elsewhere is at most N rows.
7. Appendix D.6 derives the correct cyclic reduction of $\text{tr}[X_1, X_2]^2$, but the accompanying claim that operators inside a trace commute pairwise is false and should be deleted. Cyclicity alone gives the displayed result.
8. The asymptotic formula (3.59) cannot include the endpoints $\lambda = 0, 1$ without a separate limit; at an endpoint the counting reduces to the qualitatively different one-matrix partition asymptotics. Its regime of uniformity should be stated.
9. Please replace “one restricted characters” by “one restricted character,” “when there are more than one matrix” by “when there is more than one matrix,” and use “multi-matrix” consistently.
10. Numerical tables should accompany Figure 9. They would expose the energy and coupling normalization, allow independent comparison, and make the claimed error percentages verifiable without digitizing plots.

Recommendation

The paper contains a potentially useful group-theoretic reduction of one-matrix singlet matrix elements and a promising route toward the genuinely difficult multi-matrix problem. The separation of reusable data from physical parameters is interesting. The manuscript is not yet ready for publication because its displayed one-matrix quartic decomposition is false for $N > 1$, the normalized eigenproblem is omitted, the restricted multiplicity indices are already inconsistent in (3.37), the benchmark convention is ambiguous, and the multi-matrix formula that motivates the broader physics claims has not been tested.

I recommend major revision. A successful revision should, at a minimum, replace Table 2 with a verified operator identity, regenerate all dependent matrix elements and spectra, state and verify the normalized eigenproblem, correct and test all multiplicity indices, make the Figure 9 convention reproducible, and either demonstrate (3.55) in a nontrivial two-matrix sector or narrow the paper's title and claims. Acceptance should depend on those correctness checks rather than on the present spectral plot. With those changes, a sharper novelty comparison, and corrected $U(N)/SU(N)$ physics, the work could become a valuable computational contribution.