

Referee Report

on “NLSM amplitudes from a quartic two-derivative theory”

Synopsis and overall assessment

The manuscript proposes a compact representation of planar nonlinear sigma model (NLSM) amplitudes. The authors introduce two matrix-valued fields with opposite bookkeeping polarities and the quartic two-derivative expression in equation (3). With alternating external polarities, its sole ordered vertex is asserted to be equation (4), and sums of planar quartic trees are claimed to reproduce every even-point NLSM tree amplitude. The six-point result in equation (6) is the first explicit example.

The main proof is in Appendix A. After defining the planar variables in equation (A1), the authors use factorization inductively: agreement of all lower-point residues reduces the difference between the two tree representations to a local polynomial. They then project the quartic trees onto their contact parts and organize the surviving terms by “fish” and “fishpond” configurations. Catalan recursion gives equation (A3), which matches the minimal-parametrization NLSM contact vertex in equation (A2). Generalized unitarity is then invoked to extend the result to loop integrands.

Sections III and IV develop two applications. Section III gives empirical local rules for a possible NLSM+ ϕ^3 extension, including the five quartic rules in equation (7) and the penta-vertex in equation (8), subject to global polarity restrictions. Section IV argues that the derivative placement makes complementary halves of the Adler zeros manifest, organizes the leading single-soft coefficient in equation (11), and reproduces the adjacent and next-to-adjacent double-soft terms in equations (12)–(14). The same local arguments are proposed to hold on generalized cuts at arbitrary loop order.

The tree-level idea is elegant and potentially useful. Replacing the tower of even-valent NLSM vertices by one quartic numerator gives a notably economical planar expansion, and the Catalan contact-term argument is a natural bridge to the shifted- $\text{Tr}(\Phi^3)$ and surface-kinematics literature. The six-point check and the structure of equations (A2)–(A3) support the central tree-level claim. If stated and proved at the appropriate level of precision, this would be a worthwhile contribution to the amplitudes literature.

In its present form, however, several central claims are stronger than the argument supplied. Equation (3) is not specified sufficiently to derive the ordered rules and state sums unambiguously; equality of generalized cuts is identified with equality modulo scaleless integrands without establishing the necessary completeness and regularization assumptions; and a nonrenormalizable NLSM cannot literally be an all-loop quantum theory with only its two-derivative interactions. The all-loop soft claims inherit these issues. There is also a definite combinatorial error in Table II, while the mixed extension is not yet a globally defined local theory and is used too quickly in the discussion of equation (11).

I therefore recommend **major revision**. The principal tree result looks salvageable and significant enough for a journal such as JHEP or Physical Review D, but the authors should either prove a precise bare-integrand theorem with explicit conventions or restrict the claims to the established tree-level statement and cut-level observations.

Logical structure, assumptions, and relation to prior work

The manuscript has a clear core implication:

$$\text{matching lower-point residues} + \text{matching the contact polynomial} \implies A_{2n}^\psi = A_{2n}^\pi.$$

At tree level this is the appropriate strategy. Equation (4) agrees with the four-point pion amplitude under the stated normalization, physical poles of a quartic tree separate it into lower even-point trees, and equation (A3) has the same Catalan form as equation (A2). The proof nevertheless depends on several implicit assumptions: the planar variables are treated in a specified kinematic space, the pole-free remainder has the expected two-derivative power counting, every allowed propagator cancellation is exhausted by the fishpond classification, and cyclic multiplicities and signs are those used in equation (A3). These points should be made explicit, preferably as a theorem and lemma rather than only through diagrams.

The loop statement is logically different. Tree equality implies equality of factorization data on cuts, but a conclusion about full integrands requires a definition of the integrand, loop-momentum routing, dimension, regulator, allowed numerator space, and the quotient by surface, evanescent, and cut-free terms. A conclusion about renormalized amplitudes additionally requires higher-derivative counterterms and a matching prescription. These ingredients are not consequences of the Catalan tree proof.

The paper sits close to the cited work on hidden zeros, the shifted $\text{Tr}(\Phi^3)$ representation of NLSM amplitudes, the Catalan representation of the minimal NLSM, and surface kinematics. It is also related in purpose to the known local cubic actions used for flavor–kinematics representations. The genuinely new element appears to be the particularly simple quartic, alternating-polarity organization and its fishpond proof. The manuscript should state this distinction more sharply: the Catalan contact rule and all-loop shifted-cubic containment are imported background, whereas the quartic realization and its computational economy are the proposed advance. An explicit map, or a proof that no simple map is intended, between the present rules and the shifted-cubic/surface construction would materially strengthen the novelty claim.

Major comments

1. The all-loop conclusion does not presently follow from the generalized-cut argument.

Equation (5) defines $\cong$ as equality up to scaleless integrands. Appendix A instead establishes, at most, equality on the set of generalized cuts being considered and then states that terms without cuts integrate to zero. These are not interchangeable statements without additional hypotheses. Four-dimensional cuts can miss evanescent numerator terms and rational contributions; generalized-gauge or loop-routing differences can integrate to zero without being scaleless; and a restricted set of cuts can leave contact terms with fewer propagators. Even in dimensional regularization, the manuscript must specify that the cuts are D -dimensional and complete for a stated numerator/power-counting ansatz, and must characterize the entire kernel of the cut map.

The issue is visible in the transition from the discussion following Table II to Figure 11. One one-loop example and the omission of tadpole-like graphs do not prove the assertion for arbitrary topology or loop order. Please state a precise theorem: what object is $I_n^{(L)}$, which propagator structures and external bubbles are included, what is the common loop-momentum convention, and what equivalence relation is imposed? Then either prove that every cut-free difference is a sum of integrals that vanishes in the chosen scheme, or weaken equation (5), the abstract, and the conclusion to equality of the displayed generalized cuts. A nontrivial two-loop example comparing the two global polarity choices would be especially useful, since equation (10) and Appendix A assert that their difference first becomes relevant there.

2. The quantum and renormalization scope must be separated from the bare two-derivative integrand statement.

The NLSM defined by equation (1) is an effective field theory. Loops generated by its leading two-derivative vertices are ultraviolet divergent and require higher-derivative operators with independent renormalized coefficients. The same is true of equation (3). Consequently,

the full renormalized NLSM at arbitrary loop order cannot be specified by the single quartic interaction alone. Tree-level equality of the leading-order S-matrix does not match the complete counterterm bases or their finite coefficients.

If the intended result concerns only unrenormalized integrands built from the leading-order vertices, this should be said consistently in the title-level claims, abstract, equation (5), and conclusion. If a renormalized amplitude claim is intended, the paper needs a counterterm map and a common subtraction scheme. The distinction also matters for the loop soft theorems, where local higher-derivative counterterms can modify subleading behavior even when the nonlinear symmetry preserves the leading Ward identity.

3. **Equation (3) must be defined well enough to derive equation (4).** At present the trace is said merely to record cyclic order, but equation (3) is also called a local scalar Lagrangian. These two readings are not equivalent. For ordinary matrix fields, functional differentiation of $\text{Tr}(\partial\psi^+\psi^-\partial\psi^+\psi^-)$ produces combinatorial multiplicities from the repeated ψ^+ and ψ^- insertions. The coefficient of a fixed single-trace ordering depends on the cyclic-word symmetry factor, generator normalization, and convention used in the trace decomposition. Equation (4) should therefore be derived explicitly from a component action, including all factors and the normalization of the off-diagonal propagator. The reality properties also require explanation. If ψ^+ and ψ^- are independent real fields, the off-diagonal kinetic form diagonalizes into one positive- and one negative-sign mode. If they are conjugate fields, the interaction displayed in equation (3) is not manifestly Hermitian without its conjugate. This need not be fatal if QTDS is only a formal planar diagram generator, but then the paper should avoid presenting it as an ordinary physical scalar theory and should define precisely how its diagrams and internal-state sums are generated. If it is meant as a quantum field theory, its state space, Hermiticity, unitarity status, and flavor representation must be supplied.
4. **The Catalan contact proof is compelling but presently too diagrammatic at its decisive step.** Equations (A2) and (A3) are the heart of the paper. The text asserts that, after vanishing chords are removed, each remaining blue chord corresponds to a unique propagator, that adjacent icons cannot be of the same type, and that every complete propagator cancellation is uniquely a cross with two independent fishponds. These assertions are illustrated but not proved for a general quartic tree. Please formulate the quartic tree as a planar combinatorial object and prove: (i) the precise dual-coordinate form and sign of every vertex; (ii) the one-to-one relation between cancellable numerator chords and internal edges; (iii) exhaustion and absence of overcounting of the cross/fish configurations; and (iv) the cyclic factor $1/2$ in equation (A3), including the special cases of self-complementary chords. An explicit eight-point expansion, showing both the pole residues and the full contact projection, would make the induction independently checkable. The ancillary notebook currently evaluates the QTDS expression, but it should also contain automated comparisons with an independent NLSM recursion through several multiplicities.
5. **Table II contains a definite counting error.** The column labeled “Even-valent trees” links to OEIS A003168, whose polygon-dissection interpretation gives, for $(2k + 4)$ external points,

$$4, 21, 126, 818, 5594, \dots,$$

not 4, 64, 1344, 32256, 837632. Indeed, the printed hypergeometric expression has an erroneous power of two: it evaluates to $4^k a(k)$ for the OEIS sequence $a(k)$, rather than to the required dissection number $a(k + 1)$. A convenient unambiguous formula for the intended count is

$$a(k + 1) = \frac{1}{k + 1} \sum_{j=1}^{k+1} \binom{k + 1}{j} \binom{2k + 2 + j}{j - 1}.$$

The asymptotic row must also be reindexed and renormalized consistently. The qualitative inequality with the quartic count still holds, but all numerical comparisons and any performance claim based on this column should be updated.

6. **The all-loop soft statements require a more precise cut-level analysis.** Equation (9) clearly makes one set of external momentum factors manifest at tree level, while the flipped representation makes the complementary set manifest. At loop level, however, equation (10) relies on the unproved equivalence relation discussed above. Moreover, a soft limit at fixed generic loop momentum, an algebraic surface-soft degeneration, and the soft limit of an integrated massless amplitude are distinct operations. Loop regions that scale with the soft momenta, external bubbles, and regulator effects must be addressed before a diagram-local cut argument can be promoted to an all-loop amplitude statement.

For the double-soft claim, classify all placements of the two soft legs on a general cut and show why configurations on different tree factors are subleading. For nonadjacent legs, the absence of a soft propagator alone does not make $\mathcal{O}(\tau^2)$ manifest in a fixed polarity representation when only one of the two legs carries an explicit derivative factor. One can plausibly combine analyticity with the two flipped representations, but that argument should be stated. Finally, distinguish the momentum-space limit used in equations (12)–(14) from the algebraic soft limit after Figure 11 and from the surface-soft limit in the cited all-loop literature.

7. **The derivation and notation surrounding equation (11) are not yet consistent with the mixed-sector rules.** The text writes the mixed ordering as $\phi_1\psi_2^+\dots\phi_a\dots\psi_{n-2}^+\phi_{n-1}$, apparently assigning positive polarity to every remaining ψ . Whenever a ψ block has length greater than one, this contradicts the alternating-block condition in Section III and Table I. If the superscript plus is only meant as a separator in the name of the mixed theory, it should not appear on each external field; if it is a polarity, the correct a -dependent alternating assignment must be shown.

More substantively, equation (11) is written for arbitrary multiplicity, whereas the proposed mixed rules are checked only through eleven points for a restricted class and Section III explicitly leaves their completion open. The authors should separate the established NLSM+ ϕ^3 soft theorem from a claim that the right-hand side is generated by equations (7)–(8). To retain the latter claim, prove that the relevant two-block orderings admit the stated rules at all multiplicities and show the gluing, signs, and polarity constraints explicitly. Otherwise equation (11) should be presented as an organization in terms of known mixed amplitudes, not as a consequence of the incomplete QTDS mixed extension.

8. **Section III does not yet define a finite local extension.** Equations (7) and (8) are selected ordered vertices, while the allowed polarities depend on the global external species pattern. Some patterns require a penta-vertex only after a global flip, and Appendix B identifies nine-point distributions for which the rules fail altogether. A local action would instead generate all cyclic/reflection completions and all diagrams allowed by its species content, independently of the external amplitude being targeted. The section is useful as exploratory data, but the manuscript should either give a genuine action or a complete algorithm that decides allowed vertices locally, or label the construction as a finite set of low-multiplicity identities. The abstract’s suggestion of a finite tower should be tempered unless there is evidence beyond the restricted checks. For reproducibility, the ancillary material should include the mixed-amplitude generator, the independent reference amplitudes, and machine-verifiable tests for every ordering claimed through eleven points.

Minor comments and requested clarifications

1. Equation (2) should be qualified as the tree or leading-single-trace decomposition. At loop level, multi-trace flavor structures and the precise large- N meaning of “planar” must be specified. State the

normalization of T^a and whether $\psi^\pm$ are traceless adjoint matrices.

2. In equation (5), state whether equality is meant for stripped partial amplitudes before or after summing internal flavor states. The consistency and multiplicity of polarity assignments on arbitrary cut graphs deserve a short lemma.
3. Equation (6) would be more informative if rewritten once in the planar variables of equation (A1), making its agreement with the six-point Catalan form and the two global polarity choices explicit.
4. The claimed computational improvement should be benchmarked against at least one modern NLSM recursion, not only against raw topology counts. Report the Mathematica version, hardware, whether expressions are expanded or held in factorized form, memory use, and the actual 20-point timing. The saved notebook output by itself is not a portable benchmark.
5. In equation (12), $S^{(0)}$ contains explicit τ -dependent terms. Clarify whether this is a convenient completion whose difference affects only the stated remainder, or reserve $S^{(0)}$ for the strict τ^0 limit.
6. The lower-point objects in equation (13) have one shifted, generally off-shell leg. They should be denoted as off-shell currents rather than on-shell amplitudes, and the smoothness assumption used to replace them by $A_n^\psi + \mathcal{O}(\tau)$ should be stated.
7. Figure 5 and its caption refer inconsistently to one “configuration” and plural “configurations.” More importantly, label the momentum flowing through its soft propagator so that the sign and factor in equation (14) can be checked directly.
8. Define $s_{ee'}$ and $s_{oo'}$ in the hidden-zero paragraph and say whether the statement concerns pair invariants or planar chords. Explain in which polarity representation each locus is diagram-by-diagram manifest.
9. The phrase “terms that have no generalized cuts, including scaleless integrands, integrate to zero” should be corrected even if the loop claim is weakened: only a demonstrated class of such terms may be discarded.
10. The acknowledgement should read, for example, “We thank Song He for comprehensive discussions.” A careful copy edit would also fix several comma splices, singular/plural mismatches, and the repeated four-dash heading punctuation in Appendix A.

Recommendation

The proposed quartic organization has real conceptual appeal, the tree-level contact construction appears to contain a publishable result, and the work is well aligned with the scope of JHEP or Physical Review D. Publication should nevertheless await a substantial revision that (i) defines QTDS and its ordered Feynman rules unambiguously, (ii) turns equations (A2)–(A3) into a complete general proof, (iii) corrects Table II, (iv) sharply limits or proves the all-loop and renormalized claims, and (v) disentangles the established soft results from the preliminary mixed extension. My recommendation is therefore **major revision**.