

Referee Report

Light-Front approach to $4d$ massless Higher-Spin interactions

Recommendation: Not suitable for publication in its present form. A substantially revised and more focused submission would be worth considering after the central existence and completeness claims in Chapter 4 are either proved or explicitly reduced to computational conjectures. The light-front framework, the lower-spin checks, and the quartic examples are valuable, but the present argument does not establish the advertised classification of all local quartic vertices, all unitary local higher-spin theories, or all local four-point amplitudes.

Synopsis and overall assessment

This thesis studies interactions of four-dimensional massless fields of arbitrary helicity in light-front Hamiltonian perturbation theory. Its three research chapters have a coherent progression. Chapter 2 solves the holomorphic quartic constraint, develops the “crystal” organization of products of cubic couplings, and classifies one- and two-derivative chiral theories, with and without internal color. Chapter 3 relates the same holomorphic constraint to the collinear celestial OPE, a Jacobi identity for the algebra built from cubic vertices, and the vanishing of generic chiral four-point amplitudes. Chapter 4 turns to the full, non-holomorphic quartic constraint. It constructs ansätze for quartic Hamiltonian and boost densities, reduces their determination to overdetermined first-order PDEs, tabulates candidate local solutions, proposes general helicity inequalities, identifies quasi-chiral higher-spin families, and reformulates the results in terms of factorizing spinor-helicity amplitudes.

The subject is important. Flat-space higher-spin interactions remain a sharp testing ground for locality, unitarity, and the tension between higher-spin gauge symmetry and a conventional perturbative S-matrix. The manuscript uses a technically efficient physical gauge and succeeds in recovering several nontrivial benchmarks: the Yang–Mills and gravitational quartic structures, the obstruction to interacting multiple gravitons, the known chiral higher-spin solution, and the familiar MHV amplitudes. The organization by holomorphic and anti-holomorphic derivative numbers (n, m) is useful, and the explicit higher-derivative examples suggest an interesting landscape beyond minimal couplings. The comparison between light-front constraints, celestial associativity, and amplitude factorization is also conceptually attractive.

There is, however, a decisive gap between the evidence obtained and the strength of the conclusions. The new classification in Chapter 4 rests on a finite prolongation of an overdetermined PDE system, evaluated at one kinematic point, together with finite scans in helicity and derivative number. Neither Cauchy–Kovalevskaya nor Frobenius theory gives the sufficiency claimed in Section 4.4 without an additional proof of compatibility and formal integrability. A consistent finite jet at $(x, y) = (2, 0)$ is a necessary check, not by itself an existence theorem for a local rational or analytic solution. This matters asymmetrically: a finite inconsistent subsystem can certify a no-go, but a finite consistent subsystem does not certify a yes-go. Consequently, the tables and patterns may be valuable computational evidence, but they do not presently support the words “complete,” “if and only if,” or “all.”

Several further issues amplify this concern. The reduction to the energy shell in equation (4.4.35) divides by a momentum invariant, yet the resulting vertices are later described as genuinely off shell. The amplitude argument in Section 4.7 establishes useful necessary forms from little-group scaling, dimension, and simple poles, but the converse—that every such expression factorizes into the actual cubic data of a consistent theory and that the chosen spinor basis is exhaustive—requires more proof. Finally, multiple printed formulas in the central chapters are mutually inconsistent or dimensionally incorrect. These are repairable individually, but their concentration in the formulas used to state the classification makes an independent verification particularly important.

As a doctoral thesis, the document is a substantial and generally readable synthesis. As a journal submission to a venue such as JHEP or PRD, it has an additional scope problem: Chapters 2 and 3 are identified as papers already published in JHEP, while Chapter 4 reproduces a separate submitted preprint. The publishable new unit should therefore be a focused article centered on the Chapter 4 results, with the already published material retained only as the background needed to make that article self-contained.

Logical structure and principal assumptions

The core argument can be summarized as follows.

- The physical helicity fields are used to realize the Poincare generators on the light front. Cubic Hamiltonian and boost densities are fixed by the cubic closure condition, leaving holomorphic and anti-holomorphic coupling constants labeled by three helicities.
- At quartic order, closure splits into a holomorphic constraint and a mixed constraint. The former is algebraic in products of cubic couplings; Chapters 2 and 3 classify its low-derivative solutions and connect it to a celestial OPE Jacobi identity.
- For the mixed constraint, equation (4.4.24) first restricts the numerator of the boost equation to the free-energy shell. Equation (4.4.35) then eliminates one transverse variable, and the ansatze in equations (4.4.36a)–(4.4.36h) express the local quartic density through functions $f_i(x, y)$.
- Substitution produces the coupled rational-coefficient systems (4.4.41) and (4.4.42). The manuscript prolongs these systems to a finite derivative order, evaluates them at $(x, y) = (2, 0)$, and solves the resulting algebraic equations for jets and products of cubic couplings. The output is interpreted as deciding existence or nonexistence of a local quartic density.
- Searches through $|\lambda| < 10$ for the lower-derivative tables and through $(n, m) = (8, 8)$, with some additional checks, reveal patterns. These patterns are promoted to the general conditions in Sections 4.6 and 4.7 and then used to classify unitary local theories and quasi-chiral sectors.
- A spinor-helicity ansatz, together with little-group covariance, derivative counting, and prescribed physical pole structures, produces homogeneous, single-channel, YM-like, and GR-like four-point amplitudes. The agreement of their inequalities with the light-front scan is taken as a second derivation of the general classification.

The first two bullets contain substantial established material and explicit calculations. The last four require additional arguments at precisely the points where finite computation is extrapolated to a general theorem.

Major comments

1. The PDE procedure in Section 4.4 does not presently prove existence of the asserted quartic vertices.

Equations (4.4.41) and (4.4.42) constitute an overdetermined system of first-order linear PDEs for fewer unknown functions. The paragraph below equation (4.4.42) says that Cauchy–Kovalevskaya ensures a local solution because the coefficients are analytic (indeed rational). This is not the relevant conclusion of that theorem. Cauchy–Kovalevskaya gives local existence for a system placed in an appropriate solved form with analytic, non-characteristic initial data; it does not make an arbitrary overdetermined analytic system compatible. The manuscript immediately acknowledges that compatibility conditions are required, which contradicts the preceding “always admits” statement.

The subsequent finite-jet test does not close the gap. Equation (4.4.44) represents only a finite set of differentiated relations. Footnote 22 correctly observes that checking derivatives of arbitrarily high order would guarantee integrability, but then asserts that a sufficiently high finite order ensures an actual solution. That implication needs a theorem specific to this system. One must show, for example, that the prolonged differential ideal is of finite type and has stabilized, or that the symbol is involutive and all compatibility conditions have been generated. No such rank, symbol, or differential-ideal analysis is supplied.

Evaluation at the single point $(2, 0)$ creates a second problem. Solvability of the jet equations at one point is not equivalent to a rational identity in (x, y) , and the selected point can lie on an exceptional rank locus even if no displayed denominator vanishes there. A generic-point computation must be done over the function field in x and y , or accompanied by a proof that a relevant nonzero minor stays nonzero and that all remaining equations hold on an open set.

This issue separates the two possible outputs of the calculation. If a finite subset of the exact equations is inconsistent, that is a valid certificate that the full system is inconsistent. If the subset is consistent, it merely shows that no obstruction has appeared yet. The manuscript treats both outcomes as equally decisive. The claimed yes-go entries therefore need one of the following: explicit analytic solutions h_4 and j_4 ; a proof of formal integrability and convergence; or an exact algebraic certificate that the prolongation has stabilized and defines an analytic solution. The no-go and yes-go columns should be clearly distinguished until that is done.

The computational implementation and exact input/output data should also be made available. For each table entry, it should record the rational system, denominators excluded, prolongation order, rank history, and either an explicit solution or an obstruction. This is necessary for the large tables to be independently reproducible.

2. Finite scans cannot support the advertised generic classification without an analytic induction or a revised statement of scope.

Section 4.5 states that the lower-derivative search checks helicities only up to $\lambda_{\max} = 10$. Section 4.6 says that the explicit analysis was extended through $(8, 8)$ and that the generic rules were inferred from the resulting patterns. The thesis outline is appropriately cautious when it calls the complete set a conjecture, but the abstract and later discussion alternately claim to “classify all unitary local higher-spin theories” and to “determine all local higher-spin four-point amplitudes.” Statements such as the “if and only if” following equation (4.6.9) go beyond the computation described.

The author should either provide a proof for arbitrary (n, m) and arbitrary integer helicities or consistently label the results as finite-range evidence for conjectures. A useful proof strategy would be to derive the inequalities directly from the symbol and compatibility complex of (4.4.41)–(4.4.42), then show sufficiency by constructing the corresponding density. An induction in (n, m) would also be acceptable if the map between successive ansätze and all exceptional low-order cases is made explicit. Agreement with a finite amplitude scan is a strong check, but it is not an independent proof if the amplitude argument assumes the same exhaustiveness.

The physical scope must be reduced as well. Closure through quartic order classifies quartic-order candidates, not complete interacting theories. Quintic and higher constraints may impose new relations or eliminate a spectrum. Accordingly, “fully consistent quasi-chiral theories” should be replaced by “local quartic-order solutions” unless an all-order construction is given. Likewise, a finite or infinite sum of quartic solutions is not automatically meaningful without a prescription for the spectrum, coupling convergence, and higher-point closure.

3. The passage from the on-shell reduced constraint to a genuinely off-shell Hamiltonian density needs a proof and an explicit reconstruction.

The quartic numerator is restricted to $H_2 = 0$ in equation (4.4.24), and equation (4.4.35) solves this relation by dividing by $\mathcal{P}_{12}$. The PDE system is then constructed in that patch. Later, Section 4.5 emphasizes that the resulting h_4 is a genuinely off-shell object from which one can write a Lagrangian and compute form factors. Those statements are not automatically compatible.

The author should demonstrate that every reduced solution lifts to a local off-shell representative and that the numerator in the full boost equation is divisible by H_2 , so that a local j_4 exists. Conversely, the author should show that no independent local structures proportional to H_2 , removable or non-removable by field redefinitions, are lost in the chosen ansatz. Because equation (4.4.35) divides by $\mathcal{P}_{12}$, the patch $\mathcal{P}_{12} = 0$ and analogous patches obtained by permutation require explicit treatment; a polynomial local vertex cannot be classified solely by a rational representation that is singular on such a locus.

For the Yang–Mills and gravity examples, the full off-shell h_4 and j_4 should be substituted back into the unreduced Poincare commutators as a validation. The same check should be provided for at least one genuinely new higher-derivative solution in each of the single-channel, YM-like, and GR-like classes. This would establish that the computational criterion is testing the claimed off-shell object rather than only an on-shell amplitude representative.

4. The “fast factorisation” argument in Section 4.7 does not yet establish completeness of all local four-point amplitudes.

Equations (4.7.7)–(4.7.16) use little-group weights and derivative counting to select a canonical spinor prefactor and a rational function $p(s, t)$. This efficiently produces plausible necessary forms. Equations (4.7.19)–(4.7.27) then impose simple poles and relations among their coefficients. To prove the converse claimed in the text, however, the author must show that:

- the chosen monomial times $p(s, t)$ spans all independent local four-point spinor structures modulo momentum conservation and Schouten identities, including cases in which a different bracket basis makes locality manifest;
- at each physical pole the residue equals the product of two allowed three-point amplitudes with a consistent internal helicity and the same cubic couplings that enter the other channels;
- Bose symmetry, crossing, and, where relevant, color ordering are imposed globally rather than channel by channel; and
- contact terms and amplitudes with vanishing residues are neither double counted nor omitted.

Simple physical poles plus the correct mass dimension do not by themselves guarantee factorization into the actual cubic data of a theory. The equality of the coefficients in equation (4.7.20), for example, should be derived by writing the residues with their internal-state sums, not inferred solely from the common channel form. Similarly, the manuscript should justify why the same derivative labels can be used in all permuted channels and identify the exceptional cases in which a nominal coupling vanishes.

The author already performs explicit residue matching for the YM-like example in equations (4.7.24)–(4.7.27). Extending this calculation to the generic single-channel and GR-like forms, and proving a basis theorem for the spinor structures, would turn the section into a genuine independent derivation. Until then, the result should be phrased as a classification within the stated ansatz.

5. Unitarity, parity, and the celestial claims require more precise hypotheses.

In light-front variables, parity relates holomorphic and anti-holomorphic couplings, whereas Hermitian unitarity requires the appropriate complex-conjugation relation and a positive physical-state inner product. These are not interchangeable. The discussion should use the careful distinction made in Appendix 4.A throughout Chapter 4 and should state whether all coupling solutions can be chosen to

obey the required reality conditions. A quartic local density with both chiralities is not, by itself, a proof of a unitary theory at higher orders or at the quantum level.

Chapter 3 likewise establishes a restricted and interesting equivalence: the tree-level holomorphic collinear OPE constraint, the corresponding chiral four-point relation, and a Jacobi identity coincide under the stated on-shell assumptions, with the exceptional scalar-mediated products discussed after equation (3.6.33). This should not be advertised as associativity of the full celestial CFT. The manuscript itself notes in Section 3.3 that loop OPEs contain double poles and logarithms. It also states in footnote 18 of Section 3.6 that higher-point amplitudes “should” vanish because their cuts vanish. Vanishing cuts alone does not exclude rational or contact contributions; an all-multiplicity statement needs a locality, power-counting, or recursion argument that controls such terms.

For the same reason, the assertion in Section 2.6 that vanishing tree amplitudes implies absence of one-loop UV divergences is too quick. The cited special cancellations in full chiral higher-spin gravity may use the regulated infinite spectrum and do not automatically extend to each finite spectrum classified in Chapter 2. This statement should be removed or replaced by a theory-specific loop argument, including possible rational terms and anomalies.

6. Several central displayed formulas must be corrected and the downstream calculations rechecked.

These are not merely stylistic issues:

- The functions $f_{\omega-}^{1234}$, $f_{\omega+}^{1324}$, and $f_{\omega-}^{1423}$ are already sums weighted by the exchanged helicity in equation (2.3.7). Equation (2.3.8) then multiplies them by an additional free ω outside the sum. That ω is undefined. The compact equation should presumably contain $-2C f_{\omega-}^{1234}$ and its permutations, but the author must confirm the correction and check every later use of this form.
- With the momentum parametrization of Appendix 3.B and $\beta = q^+ = (q^3 + q^0)/\sqrt{2}$, one obtains $\beta = \sqrt{2}\omega$ for an outgoing particle. Hence $\sqrt{\omega} = 2^{-1/4}\sqrt{\beta}$ up to a little-group sign or phase. Equation (3.10.9) instead gives $\sqrt{\omega} = -2^{1/4}\sqrt{\beta}$. The factor differs by $\sqrt{2}$, not merely by a sign. The normalization of equations (3.10.10)–(3.10.11) and the OPE coefficients should be checked after fixing this.
- Equation (4.6.31) begins with $\lambda_2 \leq n_1 - \lambda_2$ but concludes $\lambda_2 \leq n_2/2$. The conclusion should involve n_1 unless a different premise was intended. Equation (4.6.33) also points to n_1 .
- In the last line of equation (4.7.35), the term proportional to c'_1 contains $[12]^2[34]$ rather than $[12]^2[34]^2$. It therefore has the wrong little-group weight and mass dimension. One power of $[34]$ appears to be missing.
- Equation (4.7.37) sums the three channel terms and then adds $\mathcal{A}_s^{(2\omega-2)}$, although equation (4.7.36) already defines that quantity to include the s -channel pole. The final term should presumably be the homogeneous amplitude; otherwise the s channel is counted twice.
- The canonical definition in equation (4.7.16) is

$$d = \max\{3\lambda_i - \lambda_{jkl}, -3\lambda_i + \lambda_{jkl}\}.$$

Section 4.8 later restates the “if and only if” locality condition with $d = \max\{\lambda_i - \lambda_{jkl}, -\lambda_i + \lambda_{jkl}\}$, omitting the factors of three. Since d controls every inequality $D - d + 2k \geq 0$, this internal inconsistency must be resolved.

A symbolic dimensional, little-group, and permutation check of all displayed amplitudes and of the tables would be advisable. The corrected version should state which, if any, classifications or examples change.

Novelty, significance, and relation to the literature

The thesis sits in the light-front line initiated by Bengtsson, Bengtsson, and Brink and developed at quartic order by Metsaev. It also builds directly on the modern chiral higher-spin construction and its light-front analysis by Ponomarev, Skvortsov, and collaborators. Chapter 2’s crystal organization and low-derivative classification are useful extensions of that program, and Chapter 3 makes explicit the relation to the celestial OPE associativity and kinematic-algebra literature. These two contributions have already appeared as JHEP articles, as the manuscript states.

The genuinely prospective contribution for a new journal article is the Chapter 4 program: a systematic treatment of both boost constraints, a search for higher-derivative local quartic densities, the quasi-chiral examples, and the proposed four-class organization of amplitudes. If the existence criterion can be proved and the data made reproducible, this would be a significant result. It would sharpen the classic flat-space no-go statements by separating minimal non-abelian obstructions from higher-derivative abelian yes-go structures, and it could provide a useful bridge between off-shell light-front closure and on-shell factorization.

At present, the strongest generic conclusions remain conjectural, and the manuscript should compare them more directly with the cited work on quartic non-constructibility, local BRST cohomology, and amplitude factorization. In particular, it should explain precisely what assumption is relaxed relative to the classic higher-spin obstructions, which new structures are nontrivial modulo field redefinitions, and whether the quasi-chiral examples possess an all-order completion or are only quartic-order candidates. This would make the novelty both more credible and easier to assess.

Clarity and presentation

The broad narrative is clear, the introductory material is generous, and the many low-spin checks help orient the reader. The manuscript is nevertheless too long and repetitive for a journal article. Each research chapter retains its own introduction, notation review, conclusion, and appendices, producing substantial duplication. A focused paper should define one convention set, state the precise theorem or conjecture, present the reproducible proof or algorithm, and move extensive tables to ancillary material. The terminology should also be standardized. “Local quartic vertex,” “local four-point amplitude,” “quartic-order consistent spectrum,” and “complete local theory” are different notions and should not be used interchangeably. Similarly, “abelian vertex” refers in the cited cohomological classification to whether the gauge algebra is deformed, not to the absence of all internal tensor structure. A short paragraph fixing these meanings before Section 4.6 would prevent several ambiguities.

Minor comments

1. State at the start whether helicities are restricted to integers in every classification. Several intermediate formulas allow half-integer values, while the fermionic light-front generators and statistics are not included in the analysis.
2. Specify the allowed domain of the rational variables x and y and list all excluded divisors. The points $x = \pm 1$, $y = \pm 1$, and zeros of $\mathcal{P}_{12}$ have direct kinematic meaning and should not disappear inside a computational patch.
3. The general quartic expression in Section 4.5 writes a helicity term as $\lambda q_i / \beta_i$ inside a sum over external legs. It should be checked whether this is $\lambda_i q_i / \beta_i$.

4. Where table entries are called “all quartic vertices not listed,” append the tested helicity range in the caption. Otherwise a finite table visually reads as an unrestricted classification.
5. Give at least one explicit new higher-derivative $h_4(x, y)$ in the main text, not only relations among the coefficients k_i . This would clarify what locality means in the chosen variables.
6. In Chapter 3, distinguish the chiral celestial OPE extracted from a collinear limit from the complete celestial operator product, which may include descendants, multi-particle operators, loop singularities, and non-holomorphic terms.
7. The action discussed after the low-spin classification alternates between the notations RF^2 and FR^2 . These have different derivative and tensor content; use one defined invariant consistently.
8. Correct typographical repetitions in the ordering sums and check all section references near the end of Chapter 4; the outline refers twice to Section 4.8 where the conclusion is Section 4.9.
9. Explain whether equality of holomorphic and anti-holomorphic couplings is a parity convention or a reality condition. For complex couplings, unitarity requires complex conjugation, not simple equality.
10. The bibliography is broad, but the new Chapter 4 claims would benefit from a compact comparison table stating the assumptions and conclusions of the principal Metsaev, light-front, BRST, and on-shell no-go results cited in the text.

Required revision before reconsideration

The manuscript could become publishable if the author undertakes a focused revision with the following minimum content: (i) a mathematically valid existence criterion for the overdetermined PDEs, with explicit solutions or formal-integrability certificates for yes-go cases; (ii) reproducible exact computational data; (iii) a proof, or a clearly labeled conjecture, for the extension beyond the scanned ranges; (iv) an off-shell lifting check for the reduced solutions; (v) a complete residue and spinor-basis argument for the amplitude classification; and (vi) correction and revalidation of the formula-level inconsistencies listed above. The title, abstract, and conclusion should then distinguish established results from conjectures and quartic-order candidates from all-order theories.

The work contains enough promising structure to justify such a revision, but the central classification is not presently established to the standard required for JHEP or PRD.