

Gravity as the Square of Gauge Theory: Color–Kinematics Duality, the Double Copy, and the KLT Kernel as the Reconstruction Layer of Perturbative Spacetime

Matthew Long

The YonedaAI Collaboration

YonedaAI Research Collective

Chicago, IL

matthew@yonedaai.com · <https://yonedaai.com>

6 July 2026

Abstract

We treat perturbative gravity amplitudes not as output of the Einstein–Hilbert action expanded into an infinite tower of vertices, but as reconstructions of gauge-theory data. The mechanism is the double copy of Bern, Carrasco and Johansson: once one exhibits a representation of a gauge-theory amplitude in which the kinematic numerators obey the same Jacobi and antisymmetry relations as the color factors—color–kinematics duality—the corresponding gravity amplitude is obtained by discarding the color factors and squaring the kinematic numerators. In this reading the gauge-theory kinematic numerator is the invariant representation datum, and gravity is its algebraic image; no gravitational input enters beyond the choice of double-copy pairing. This is the amplitude-level face of the project’s governing thesis, that spacetime and its scattering data are reconstructed from invariant representation structures rather than posited.

We give a self-contained account of the cubic-graph (trivalent) organization of amplitudes, the color Jacobi relations and their kinematic mirror, the double-copy substitution $c_g \mapsto \tilde{n}_g$, and the Kawai–Lewellen–Tye (KLT) relations that historically preceded it. Our statements are labelled propositions and each carries the project’s S/H/P epistemic tag (standard / heuristic / speculative). The BCJ amplitude relations, the existence of dual numerators at tree level, the tree double copy, and the identification of the field-theory KLT kernel with the inverse of the biadjoint-scalar amplitude matrix are standard amplitude mathematics [S]; their reading as statements about a physical gravitational S-matrix is heuristic and confined to the double-copy-constructible theories [S/H]; the loop-level double copy and the classical (Kerr–Schild, Weyl) double copy are heuristic in their established sectors [S/H] and speculative as general claims [H]; and coupling the numerator to the motivic-coaction “residue tree = coaction tree” conjecture caps the joint claim at [H/P]. We work the four-graviton amplitude in full:

an explicit color-dual numerator set, the double-copy assembly, the four-point KLT relation, and the identification of the KLT kernel with the inverse biadjoint amplitude. We sketch the five-point structure, and we treat the classical double copy through Kerr–Schild and Weyl. We locate the regime among its neighbours—the associahedron canonical form of the companion positive-geometry paper, the motivic coaction, and the celestial change of basis—tracking epistemic status under the worst-component-wins composition calculus. We close with honest obstructions: the absence of a general all-loop existence proof of dual numerators, the dilaton/axion contamination of the naive square, and the confinement of the classical double copy to algebraically special backgrounds. A companion Haskell development builds the color factors of the four-point amplitude, verifies the color Jacobi identity and its kinematic mirror exactly over the rationals, assembles the double copy, and checks crossing symmetry and generalized gauge invariance.

Contents

1	Introduction	3
1.1	The reconstruction thesis in the amplitude regime	3
1.2	What is and is not claimed	4
1.3	Relation to the companion papers and the roadmap	4
1.4	Organization	4
2	The representation dictionary and the S/H/P calculus	5
2.1	The representation stack and realization pipeline	5
2.2	Warrants and worst-component composition	5
2.3	The dictionary entry for this paper	6
3	Cubic-graph organization and the color algebra	6
3.1	Trivalent graphs and the color decomposition	6
3.2	Color factors and the Jacobi identity	7
3.3	Color-ordered partial amplitudes	7
4	Color–kinematics duality	8
4.1	Definition	8
4.2	Existence at tree level	8
4.3	The kinematic algebra	9
5	The double copy	9
5.1	The substitution rule	9
5.2	The spectrum of the square	9
5.3	Tree-level correctness and equivalence to KLT	10
6	The KLT kernel as the inverse biadjoint amplitude	11
6.1	The biadjoint scalar	11

7	Worked example: the four-point amplitude	12
7.1	Kinematics	12
7.2	Color factors	12
7.3	Color-dual numerators	12
7.4	The double copy	13
7.5	Matching KLT	13
7.6	Five points, sketched	14
8	The classical double copy	15
8.1	The Kerr–Schild double copy	15
8.2	The Weyl double copy	15
8.3	Reconstruction reading	16
9	Relation to neighbouring regimes	16
9.1	Positive geometry (companion QG-IV paper)	16
9.2	Motivic coaction and open problem 1	16
9.3	Cut/regulator complementarity and open problem 2	17
9.4	Celestial holography	17
9.5	Summary table	17
10	Honest limitations and open problems	17
11	The companion Haskell development	18
12	Conclusion	19

1 Introduction

1.1 The reconstruction thesis in the amplitude regime

The governing perspective of this project is that quantum gravity is not, in the first instance, the quantization of material objects placed in a pre-existing spacetime. It is the reconstruction of spacetime, matter, causality, and measurement from invariant representation structures. Different regimes of the theory instantiate the idea with different invariants. The present paper concerns perturbative gravity amplitudes, where the reconstruction is unusually literal: the gravitational S-matrix, order by order in the coupling, is an algebraic square of gauge-theory data, and the datum that gets squared is a representation-theoretic object—a kinematic numerator obeying a color-like algebra.

The conventional route to a gravity amplitude begins with the Einstein–Hilbert action, expands $g_{\mu\nu} = \eta_{\mu\nu} + \kappa h_{\mu\nu}$, and generates an infinite series of increasingly complicated vertices. Even the three-graviton vertex has of order a hundred terms; the four-point tree amplitude assembled directly from Feynman diagrams is a notorious exercise. The theory looks, from this vantage, both intractable and non-renormalizable. The double copy inverts the situation. It asserts that the same amplitude is the square of a Yang–Mills amplitude, and Yang–Mills amplitudes are comparatively simple. The technical content that makes the squaring possible

is color–kinematics duality: a Yang–Mills amplitude can be written so that its kinematic numerators satisfy the same Jacobi identities as the Lie-algebra color factors that accompany them. Once this is arranged, gravity is obtained by a substitution rule—delete the color, put a second numerator in its place—with no independent gravitational input.

This is a sharp instance of reconstruction from representation structure. The invariant is not the gravitational field. It is the kinematic numerator, an object living in the same algebraic world as the adjoint-representation color factor of a gauge theory. Gravity is what one obtains by pairing two such invariants. The classical counterpart is equally striking: exact black-hole and gravitational-wave solutions of general relativity are squares of exact solutions of Maxwell or Yang–Mills theory, through the Kerr–Schild and Weyl double copies. Curvature, like the S-matrix, is reconstructed rather than posited.

1.2 What is and is not claimed

We are careful throughout to separate mathematics from physics and both from speculation, using the project’s three-valued epistemic calculus (Section 2). The tree-level statements we prove or cite—the BCJ amplitude relations, the existence of color-dual numerators at tree level, the tree double copy, the KLT relations, and the identification of the field-theory KLT kernel with the inverse of the biadjoint-scalar partial-amplitude matrix—are standard mathematics of scattering amplitudes, tag [S]. Read as statements about a physical gravitational S-matrix they become heuristic, tag [S/H], because they hold for the specific spectrum of double-copy-constructible theories (pure Yang–Mills squared gives $\mathcal{N} = 0$ supergravity, not pure Einstein gravity; $\mathcal{N} = 4$ super-Yang–Mills squared gives $\mathcal{N} = 8$ supergravity) rather than for a general ultraviolet-complete theory of quantum gravity. The loop-level double copy is verified in a large number of explicit high-loop cases but lacks a general existence proof and is tagged [H]. The classical double copy is [S/H] in its Kerr–Schild and Weyl sectors and [H] as a general claim. The conjectural bridge to motivic coactions, shared with the companion positive-geometry paper, is [H/P]. We never silently upgrade a label.

1.3 Relation to the companion papers and the roadmap

This paper is the double-copy half of module QG-IV (amplitudes, positive geometry, double copy) in the project roadmap. Its companion, the positive-geometry paper, treats the same module through canonical forms of positive geometries; the two meet at a single sharp object, the KLT kernel, which in the positive-geometry paper is the inverse of the canonical form of the Arkani-Hamed–Bai–He–Yan associahedron (the biadjoint-scalar amplitude) and in this paper is the bilinear form pairing two gauge-theory copies into gravity. That these are the same object (Theorem 6.2) is the technical seam between the two papers. The regime also touches open problems 1 (residue tree = coaction tree) and 2 (cut/regulator complementarity) of the roadmap queue, which we discuss honestly as open in Sections 9 and 10.

1.4 Organization

Section 2 recalls the representation-stack framework and the S/H/P calculus. Section 3 sets up the cubic-graph organization of amplitudes and the color algebra. Section 4 defines

color–kinematics duality and the kinematic Jacobi relations. Section 5 states and proves the tree double copy and its equivalence to KLT. Section 6 identifies the KLT kernel with the inverse biadjoint amplitude, the seam to the positive-geometry paper. Section 7 works the four-point example in full and sketches five points. Section 8 treats the classical Kerr–Schild and Weyl double copies. Section 9 locates the regime among its neighbours under the weakest-link calculus. Section 10 states the honest limitations and open problems. Section 11 describes the companion Haskell development. Section 12 concludes.

2 The representation dictionary and the S/H/P calculus

We recall the minimal apparatus this paper shares with the other eleven topic papers of the library, so that the paper is self-contained.

2.1 The representation stack and realization pipeline

A representation entry is a tuple $E = (M, P, \tau, \sigma)$: a mathematical object M , a physical target P , a translation datum τ that carries the correspondence together with its assumptions and limitations, and an epistemic status σ . The realization pipeline

$$M \longrightarrow \Phi(M) \longrightarrow \text{Real}_\alpha(\Phi(M)) \longrightarrow \text{physical representation} \longrightarrow \text{Obs}$$

sends a mathematical object to observables through a choice of realization α . In the present regime M is the system of color-dual kinematic numerators of a gauge theory, Real is the double-copy substitution, and Obs is a gravity amplitude.

2.2 Warrants and worst-component composition

There are three ordered warrants,

$$S < H < P,$$

read as: S , a standard mathematical or mathematical-physics correspondence; H , a strong heuristic physical representation; P , a speculative ontological extension. A status is a nonempty tuple of warrants, rendered $\mathbf{S}$, $\mathbf{H}$, $\mathbf{P}$, or a composite such as $\mathbf{S/H}$. The single structural rule we use repeatedly is that composition is monotone and the worst component wins.

Definition 2.1 (Status composition). For warrants $a, b \in \{S, H, P\}$ define $a \vee b = \max(a, b)$ under $S < H < P$. For statuses (tuples), the composite is the join of worst components: $\sigma_1 \circ \sigma_2$ has worst component $\max(\text{worst}(\sigma_1), \text{worst}(\sigma_2))$.

Proposition 2.2 (Monotonicity and unit). S is a unit for $\vee$: $S \vee a = a$. The operation is associative, commutative, and idempotent, and $\text{worst}(\sigma_1 \circ \sigma_2) \geq \text{worst}(\sigma_i)$ for each i . In particular no chain of representations has a composite status better than its weakest link. [S]

Proof. $(\{S, H, P\}, \vee)$ is the join-semilattice of a totally ordered three-element set; $\max$ is associative, commutative, idempotent, and has the least element S as unit. The inequality is monotonicity of $\max$. These are elementary lattice facts. $\square$

This proposition is the engine of Section 9: chaining the double copy to the motivic coaction, or to celestial holography, cannot yield a composite more secure than the H -rated conjectural links in those chains.

2.3 The dictionary entry for this paper

The relevant row of the library’s dictionary is

Status	Math	Physical representation
S/H	Color–kinematics dual numerator system	Gravity amplitude as a gauge-theory double copy

with assumptions “a suitable color–kinematics representation exists” and limitations “loop-level and theory-specific constructions require care.” Everything below refines this single entry.

3 Cubic-graph organization and the color algebra

3.1 Trivalent graphs and the color decomposition

Fix a gauge group with Lie algebra $\mathfrak{g}$, structure constants f^{abc} defined by $[T^a, T^b] = i f^{abc} T^c$ in a basis of Hermitian generators normalized to $\text{Tr}(T^a T^b) = \delta^{ab}$. Consider the L -loop n -point amplitude of a gauge theory whose fields are all in the adjoint representation. The central organizational fact is that every Feynman diagram can be reduced to trivalent (cubic) topology: quartic vertices are rewritten by inserting and cancelling a propagator, so that contact terms are absorbed into cubic graphs at the price of numerators that are polynomials in the momenta.

Definition 3.1 (Cubic-graph representation). Let $\Gamma_{n,L}$ be the set of distinct trivalent graphs with n external legs and L loops (no tadpoles or external self-energies). A *cubic-graph representation* of the L -loop n -point amplitude is an expression

$$\mathcal{A}_n^{(L)} = i^L g^{n-2+2L} \sum_{g \in \Gamma_{n,L}} \frac{1}{S_g} \int \prod_{\ell=1}^L \frac{d^D p_\ell}{(2\pi)^D} \frac{c_g n_g}{\prod_{\alpha \in g} D_\alpha}, \quad (1)$$

where D_α are the inverse propagators (squared momenta) of the internal lines of g , c_g is the color factor obtained by dressing each trivalent vertex with a factor f^{abc} and contracting along the graph, n_g is a kinematic numerator (a polynomial in momenta and polarizations), and S_g is the symmetry factor of g .

At tree level, $L = 0$, the momentum integrals and symmetry factors are trivial and (1) reads

$$\mathcal{A}_n^{\text{tree}} = g^{n-2} \sum_{g \in \Gamma_{n,0}} \frac{c_g n_g}{\prod_{\alpha \in g} D_\alpha}. \quad (2)$$

The number of cubic tree graphs with n external legs is $(2n-5)!! = 1 \cdot 3 \cdot 5 \cdots (2n-5)$: three graphs at $n = 4$, fifteen at $n = 5$, and one hundred and five at $n = 6$.

3.2 Color factors and the Jacobi identity

The color factors c_g are not independent. Whenever three cubic graphs g_s, g_t, g_u differ only by the connectivity of a single internal four-point subgraph—an s -, t -, or u -channel exchange in one propagator slot, with everything else held fixed—their color factors satisfy a linear relation inherited from the Lie-algebra Jacobi identity.

Lemma 3.2 (Color Jacobi relation). *Let g_s, g_t, g_u be three cubic graphs identical outside a distinguished four-point subgraph, in which they realize the three inequivalent channels. With the standard sign convention on edge orientations,*

$$c_{g_s} + c_{g_t} + c_{g_u} = 0. \quad (3)$$

Moreover $c_g \mapsto -c_g$ under the flip of any single internal edge (antisymmetry). [S]

Proof. The three channels differ only in the way four adjacent structure constants are contracted across the distinguished subgraph. Writing the four external adjoint indices of the subgraph as a, b, c, d and the shared internal index as e , the three contractions are $f^{abe} f^{ecd}$, $f^{bce} f^{ead}$, and $f^{cae} f^{ebd}$. The Jacobi identity of $\mathfrak{g}$, $f^{abe} f^{ecd} + f^{bce} f^{ead} + f^{cae} f^{ebd} = 0$, gives (3). Antisymmetry of f^{abc} in its indices gives the sign flip. $\square$

Example 3.3 (Four points). At $n = 4$ there are exactly three cubic graphs, the s -, t -, and u -channel exchanges, and $\Gamma_{4,0} = \{g_s, g_t, g_u\}$. Their color factors are

$$c_s = f^{a_1 a_2 e} f^{e a_3 a_4}, \quad c_t = f^{a_2 a_3 e} f^{e a_1 a_4}, \quad c_u = f^{a_3 a_1 e} f^{e a_2 a_4},$$

and (3) is the single relation $c_s + c_t + c_u = 0$, the Jacobi identity itself. The companion Haskell code verifies this over $\mathfrak{su}(2)$ (Section 11), where $f^{abc} = \varepsilon^{abc}$ and the relation is checked as an exact identity of integer tensors.

3.3 Color-ordered partial amplitudes

Contracting (2) against a trace basis organizes the amplitude into color-ordered partial amplitudes. At tree level, for adjoint fields,

$$\mathcal{A}_n^{\text{tree}} = g^{n-2} \sum_{\sigma \in S_n / \mathbb{Z}_n} \text{Tr}(T^{a_{\sigma(1)}} \cdots T^{a_{\sigma(n)}}) A(\sigma(1), \dots, \sigma(n)), \quad (4)$$

where the partial amplitudes $A(1, \dots, n)$ are gauge-invariant, cyclically symmetric, and carry only the kinematic data. The $(n-1)!$ orderings in (4) are not independent: the Kleiss–Kuijf relations reduce them to a basis of $(n-2)!$, and, as we recall next, color–kinematics duality reduces that further to $(n-3)!$.

4 Color–kinematics duality

4.1 Definition

The color factors obey the Jacobi relations (3) by Lie algebra. There is no a priori reason for the kinematic numerators to obey anything of the sort; they are, in a generic gauge and diagram-by-diagram, arbitrary. The discovery of Bern, Carrasco and Johansson is that the gauge freedom in (1)—the freedom to move contact terms between graphs—is exactly enough to force the numerators into the same algebra as the colors.

Definition 4.1 (Color–kinematics duality). A cubic-graph representation (1) is *color–kinematics dual* (or the numerators are *BCJ dual*, or *color–dual*) if the kinematic numerators satisfy the same linear and sign relations as the color factors: for every Jacobi triple g_s, g_t, g_u ,

$$c_{g_s} + c_{g_t} + c_{g_u} = 0 \implies n_{g_s} + n_{g_t} + n_{g_u} = 0, \quad (5)$$

and $n_g \mapsto -n_g$ under the flip of any single internal edge whenever $c_g \mapsto -c_g$.

Remark 4.2. The condition (5) is a condition on the *numerators*, not on the amplitude. The amplitude (1) is invariant under *generalized gauge transformations* $n_g \mapsto n_g + \Delta_g$ with $\sum_g \Delta_g c_g / \prod D = 0$; color–kinematics duality is the statement that some representative in this orbit satisfies (5). That such a representative exists is nontrivial and is the content of the next subsection.

4.2 Existence at tree level

Theorem 4.3 (Existence of dual numerators at tree level). *For every tree-level n -point amplitude of Yang–Mills theory (and of a large class of adjoint-coupled theories) there exists a color–kinematics dual cubic-graph representation. Equivalently, the partial amplitudes satisfy the fundamental BCJ relations, which reduce the $(n - 2)!$ -dimensional Kleiss–Kuijf basis to an $(n - 3)!$ -dimensional minimal basis. [S]*

Proof sketch. Two independent proofs are known. The first, due to Bern, Carrasco and Johansson and completed via on-shell recursion by Feng, Huang and Jia and by others, derives the fundamental BCJ relations directly from BCFW recursion on the partial amplitudes; the existence of dual numerators is then equivalent to the solvability of the linear system relating numerators to a minimal basis of partial amplitudes, which the BCJ relations render consistent. The second, due to Cachazo, He and Yuan, is constructive: the scattering equations localize the tree amplitude on the moduli space $\mathcal{M}_{0,n}$ of n marked points on the sphere, and the resulting Cachazo–He–Yuan representation manifestly satisfies both the Kleiss–Kuijf and the fundamental BCJ relations, so the numerators read off from it are dual. The monodromy relations of string theory provide a third route: the fundamental BCJ relations are the field-theory limit of the monodromy relations among open-string partial amplitudes, which follow from contour deformation of the Koba–Nielsen integral. We take the $(n - 3)!$ count as the operative statement; the four-point instance ($1! = 1$) is worked explicitly in Section 7. $\square$

Remark 4.4 (Existence versus locality). At tree level dual numerators can always be chosen local (polynomial in momenta). At loop level this is no longer guaranteed in general; the tension between manifest duality and manifest locality is one of the central technical difficulties of the loop-level program (Section 10).

4.3 The kinematic algebra

Color–kinematics duality suggests that the numerators are structure constants of a “kinematic algebra” dual to the gauge Lie algebra. In the self-dual sector of Yang–Mills this algebra was identified by Monteiro and O’Connell as the algebra of area-preserving diffeomorphisms; for the full theory a closed-form kinematic algebra is known in specific sectors. Brandhuber, Chen, Johansson, Travaglini and Wen gave closed-form BCJ numerators for heavy-mass effective field theory and Yang–Mills through a quasi-shuffle Hopf algebra, making the duality manifest in those sectors. A general, all-multiplicity, all-loop kinematic algebra is not known; we return to this in Section 10. The status of the kinematic-algebra program as physics is [H]: it is a strong structural hypothesis with explicit realizations, not a theorem in generality.

5 The double copy

5.1 The substitution rule

Definition 5.1 (Double copy). Given a color–kinematics dual representation (1) of a gauge-theory amplitude with numerators n_g , and a second (not necessarily dual) representation of a gauge theory sharing the same graphs and propagators with numerators $\tilde{n}_g$, the associated *double-copy* amplitude is

$$\mathcal{M}_n^{(L)} = i^L \left(\frac{\kappa}{2}\right)^{n-2+2L} \sum_{g \in \Gamma_{n,L}} \frac{1}{S_g} \int \prod_{\ell=1}^L \frac{d^D p_\ell}{(2\pi)^D} \frac{n_g \tilde{n}_g}{\prod_{\alpha \in g} D_\alpha}. \quad (6)$$

The color factors of (1) are replaced by a second copy of kinematic numerators; the propagators are unchanged.

Remark 5.2 (Only one copy need be dual). It suffices that one of the two numerator sets satisfy (5). If n_g is dual, then (6) is invariant under generalized gauge transformations of the $\tilde{n}_g$, so the second set may be taken in any convenient gauge. This asymmetry is important in practice: one squares a hard, color-dual numerator against an easy, generic one.

5.2 The spectrum of the square

The field content of the double copy is the tensor product of the two gauge-theory field contents. A gauge boson has $D - 2$ physical polarizations; the product of two gauge bosons decomposes under the little group as

$$(D - 2) \otimes (D - 2) = \underbrace{\text{Sym}_0^2}_{\text{graviton}} \oplus \underbrace{\wedge^2}_{B\text{-field}} \oplus \underbrace{\text{trace}}_{\text{dilaton}}.$$

Thus pure Yang–Mills squared is *not* pure Einstein gravity: it is $\mathcal{N} = 0$ supergravity, the graviton accompanied by a two-form (Kalb–Ramond) field and a scalar dilaton (in four dimensions the two-form is dual to a pseudoscalar axion, which is the “axion” of the abstract). Projecting onto pure gravity requires removing the dilaton and two-form, which is done by adding ghost matter in the double copy (Johansson–Ochirov) or by a suitable projection. Squaring $\mathcal{N} = 4$ super-Yang–Mills gives $\mathcal{N} = 8$ supergravity. We flag this explicitly: the slogan “gravity = gauge²” is exact only with the correct bookkeeping of the extra states.

5.3 Tree-level correctness and equivalence to KLT

Theorem 5.3 (Tree-level double copy). *Let n_g be a color–kinematics dual tree-level numerator set for a gauge theory, and $\tilde{n}_g$ a tree-level numerator set for a second gauge theory sharing the same cubic graphs. Then (6) at $L = 0$ is the tree amplitude of the corresponding (super)gravity theory, and it is independent of the generalized gauge chosen for $\tilde{n}_g$. Moreover the resulting expression coincides with the Kawai–Lewellen–Tye form*

$$\mathcal{M}_n^{\text{tree}} = \sum_{\sigma, \tau} A(\sigma) \mathcal{S}[\sigma|\tau] \tilde{A}(\tau), \quad (7)$$

where σ, τ run over an $(n-3)!$ minimal basis of orderings and $\mathcal{S}[\sigma|\tau] = (m^{-1})[\sigma|\tau]$ is the field-theory KLT kernel, the inverse of the biadjoint amplitude matrix of Theorem 6.2. All ordering-dependent signs are carried by the kernel, so (7) agrees with the matrix form $\mathcal{M}_n = A m^{-1} \tilde{A}^\top$ of Theorem 6.2 with no extra prefactor. Here $A(\sigma), \tilde{A}(\tau)$ denote the coupling-stripped partial amplitudes (the g^{n-2} prefactor of (2) removed); the coupling bookkeeping is carried by the gravitational prefactor $(\kappa/2)^{n-2}$ of (6), so the replacement $g^{n-2} g^{n-2} \rightarrow (\kappa/2)^{n-2}$ is part of the double-copy map and not an equality of prefactors. [S] as amplitude mathematics, [S/H] read as the physical gravity S-matrix.

Proof sketch. Gauge independence of (6) follows from duality of n_g : a generalized gauge shift $\tilde{n}_g \mapsto \tilde{n}_g + \tilde{\Delta}_g$ changes (6) by $\sum_g n_g \tilde{\Delta}_g / \prod D$, which vanishes because the n_g satisfy the same Jacobi relations as color factors and $\tilde{\Delta}_g$ is by definition a combination that annihilates against color, hence against any dual set. Equivalence to KLT is proved by expressing both (6) and (7) in the minimal $(n-3)!$ basis of partial amplitudes: the numerator bilinear reorganizes, using the fundamental BCJ relations of Theorem 4.3, into the bilinear (7) with kernel the inverse of the biadjoint amplitude matrix (Theorem 6.2). Correctness against the direct gravity amplitude was established by Bern, Carrasco and Johansson for n points at tree level, and independently follows from the Cachazo–He–Yuan formula, whose gravity version is manifestly the square of its gauge version on $\mathcal{M}_{0,n}$. $\square$

Corollary 5.4. *At $n = 4$ the minimal basis is one-dimensional, the kernel is the single number $\mathcal{S} = m^{-1} = -s_{12}$ (Theorem 7.2), and (7) reduces to*

$$\mathcal{M}_4^{\text{tree}} = -s_{12} A(1, 2, 3, 4) \tilde{A}(1, 2, 4, 3), \quad (8)$$

with $s_{12} = (k_1 + k_2)^2$, where $A, \tilde{A}$ are the physical color-ordered partial amplitudes of the two gauge copies, whose relative channel signs are fixed by the trace decomposition. [S]

6 The KLT kernel as the inverse biadjoint amplitude

The KLT kernel is the one object shared with the companion positive-geometry paper. There it is the inverse of a canonical form; here it is the pairing that makes gravity a square. The identification is a theorem.

6.1 The biadjoint scalar

Definition 6.1 (Biadjoint scalar amplitude). Let ϕ^3 biadjoint scalar theory be the theory of a scalar $\phi^{a\bar{a}}$ in the adjoint of $\mathfrak{g} \times \tilde{\mathfrak{g}}$ with cubic interaction $f^{abc} \tilde{f}^{\bar{a}\bar{b}\bar{c}} \phi^{a\bar{a}} \phi^{b\bar{b}} \phi^{c\bar{c}}$. Its tree amplitudes are doubly color-ordered; stripping both orderings gives the *doubly-partial amplitudes* $m(\sigma|\tau)$, sums of products of scalar propagators over the trivalent graphs compatible with both orderings σ and τ , with signs fixed by the planar orderings.

The matrix $m(\sigma|\tau)$, indexed by an $(n-3)!$ minimal basis of orderings, is the central combinatorial object of the biadjoint theory. In the companion paper it is the canonical form of the Arkani-Hamed–Bai–He–Yan associahedron.

Theorem 6.2 (KLT kernel = inverse biadjoint amplitude). *The field-theory KLT kernel $\mathcal{S}[\sigma|\tau]$ of (7), restricted to an $(n-3)!$ minimal basis, is the inverse of the biadjoint doubly-partial-amplitude matrix $m(\sigma|\tau)$ on that basis:*

$$\mathcal{S} = m^{-1}, \quad \mathcal{M}_n^{\text{tree}} = A m^{-1} \tilde{A}^\top, \quad (9)$$

where $A, \tilde{A}$ are the row vectors of minimal-basis partial amplitudes of the two gauge copies. [S]

Proof sketch. Two derivations are standard. Cachazo, He and Yuan observed that the scattering-equations integrand for gravity is the square of that for gauge theory, and that the biadjoint scalar sits between them as the object whose doubly-ordered amplitudes assemble the change of basis; inverting the resulting m -matrix on the minimal basis yields the kernel. Mizera gave a proof via intersection theory of twisted cocycles on $\mathcal{M}_{0,n}$: the KLT kernel is the intersection pairing of twisted cycles, the biadjoint amplitude is the intersection pairing of the dual cocycles, and the two pairings are inverse by the general duality of twisted homology and cohomology. Either way $\mathcal{S} = m^{-1}$. The four-point instance is Theorem 7.2. $\square$

Remark 6.3 (The seam to the companion paper). Theorem 6.2 is the technical junction between this paper and the positive-geometry paper. There m is produced as the canonical form of the associahedron—a top-degree differential form whose logarithmic singularities sit on the boundary strata of a polytope. Here m^{-1} is produced as the bilinear pairing that squares two gauge theories into gravity. That the amplitude combinatorics of a polytope and the algebra of the double copy meet in one invertible matrix is, in the language of the project, the statement that the two QG-IV representations realize the same invariant.

7 Worked example: the four-point amplitude

We now carry the entire construction through explicitly at $n = 4$, where every object is a finite rational function of the Mandelstam invariants and every claim is checkable by hand and in code.

7.1 Kinematics

For four massless momenta $k_1, \dots, k_4$ with $\sum_i k_i = 0$ and $k_i^2 = 0$, define

$$s = s_{12} = (k_1 + k_2)^2, \quad t = s_{23} = (k_2 + k_3)^2, \quad u = s_{13} = (k_1 + k_3)^2.$$

Momentum conservation and masslessness give the single constraint

$$s + t + u = 2(k_1 \cdot k_2 + k_2 \cdot k_3 + k_1 \cdot k_3) = (k_1 + k_2 + k_3)^2 = k_4^2 = 0, \quad (10)$$

using $k_i^2 = 0$ and $k_1 + k_2 + k_3 = -k_4$. This is the only relation among the three invariants and it is exact.

7.2 Color factors

The three cubic graphs g_s, g_t, g_u carry the color factors of Theorem 3.3, and by Theorem 3.2,

$$c_s + c_t + c_u = 0. \quad (11)$$

Over $\mathfrak{su}(2)$ with $f^{abc} = \varepsilon^{abc}$ this is the identity $\varepsilon^{a_1 a_2 e} \varepsilon^{e a_3 a_4} + \varepsilon^{a_2 a_3 e} \varepsilon^{e a_1 a_4} + \varepsilon^{a_3 a_1 e} \varepsilon^{e a_2 a_4} = 0$, which expands via $\varepsilon^{abe} \varepsilon^{cde} = \delta^{ac} \delta^{bd} - \delta^{ad} \delta^{bc}$ into a sum of six Kronecker-delta products that cancel in pairs; the companion code checks it for all 3^4 index assignments.

7.3 Color-dual numerators

The color-ordered partial amplitude in the ordering $(1, 2, 3, 4)$ receives contributions from the s - and t -channels only, and the two channels enter with a relative sign fixed by the trace decomposition (the color factors c_s and c_t dress the traces $\text{Tr}(1234)$ and $\text{Tr}(1243)$ with opposite sign, which arises mechanically from expanding each structure constant as a commutator, $f^{abc} \propto \text{Tr}([T^a, T^b]T^c)$, so the two orderings inside each vertex enter with opposite sign). With that convention,

$$A(1, 2, 3, 4) = \frac{n_s}{s} - \frac{n_t}{t}, \quad A(1, 2, 4, 3) = -\frac{n_s}{s} + \frac{n_u}{u}, \quad (12)$$

the second following from the same rule applied to the ordering $(1, 2, 4, 3)$, whose planar channels are s and u . These signs are structural, not cosmetic: they are exactly what makes the KLT bilinear (8) reproduce the double-copy sum, as we verify below. A numerator set is color-dual when

$$n_s + n_t + n_u = 0, \quad (13)$$

mirroring (11). To make the arithmetic completely explicit we use a toy dual set that solves (13) identically,

$$n_s = t - u, \quad n_t = u - s, \quad n_u = s - t, \quad (14)$$

so that $n_s + n_t + n_u = (t - u) + (u - s) + (s - t) = 0$ with no use of (10). These are the numerators of a scalar theory whose partial amplitude is $A(1, 2, 3, 4) = (t - u)/s + (u - s)/t$; they are genuinely color-dual and suffice to exhibit every structural feature of the double copy. (For the physical gluon amplitude the numerators are the corresponding polarization-dependent expressions; the algebra of the double copy is identical and we keep the scalar toy for transparency.)

Remark 7.1 (Generalized gauge freedom, explicitly). Under the uniform shift $n_g \mapsto n_g + D_g \alpha$ (same α on all three channels) the partial amplitude (12) is invariant: with the relative trace sign, $A(1, 2, 3, 4) \mapsto (n_s + s\alpha)/s - (n_t + t\alpha)/t = A(1, 2, 3, 4) + \alpha - \alpha = A(1, 2, 3, 4)$. The full amplitude (2) is invariant under the same shift because $\sum_g D_g \alpha c_g / D_g = \alpha \sum_g c_g = \alpha(c_s + c_t + c_u) = 0$ by (11). Thus color Jacobi is precisely what makes generalized gauge transformations invisible in the amplitude; the companion code checks this shift invariance of the full color-dressed amplitude numerically over the rationals.

7.4 The double copy

Applying (6) at $n = 4$, $L = 0$ with $\tilde{n}_g = n_g$ (squaring one theory against itself),

$$\mathcal{M}_4 = \frac{n_s^2}{s} + \frac{n_t^2}{t} + \frac{n_u^2}{u} = \frac{(t - u)^2}{s} + \frac{(u - s)^2}{t} + \frac{(s - t)^2}{u}. \quad (15)$$

This expression is totally symmetric under permutations of (s, t, u) accompanied by the induced permutation of the numerators, hence crossing symmetric, as a gravity amplitude must be. Using $u = -s - t$ one can reduce (15) to a rational function of s, t alone; the companion code evaluates both (15) and its reduced form on random rational kinematics and checks they agree, and checks the crossing symmetry directly.

7.5 Matching KLT

By Theorem 5.4, the same amplitude is the KLT bilinear

$$\mathcal{M}_4 = -s A(1, 2, 3, 4) \tilde{A}(1, 2, 4, 3), \quad (16)$$

with $A, \tilde{A}$ the physical color-ordered partial amplitudes (12) of the two copies. Taking both copies equal to the toy dual set (14), so that $A(1, 2, 3, 4) = n_s/s - n_t/t$ and $\tilde{A}(1, 2, 4, 3) = -n_s/s + n_u/u$, expand (16):

$$\begin{aligned} -s A(1, 2, 3, 4) \tilde{A}(1, 2, 4, 3) &= -s \left(\frac{n_s}{s} - \frac{n_t}{t} \right) \left(-\frac{n_s}{s} + \frac{n_u}{u} \right) \\ &= \frac{n_s^2}{s} - \frac{n_s n_u}{u} - \frac{n_t n_s}{t} + \frac{s n_t n_u}{tu}. \end{aligned} \quad (17)$$

Substituting $n_s = -(n_t + n_u)$ in the last three terms and collecting over the common denominator tu gives $n_t n_u (s + t + u)/(tu) + n_t^2/t + n_u^2/u$; the first term vanishes by (10), leaving

$$-s A(1, 2, 3, 4) \tilde{A}(1, 2, 4, 3) = \frac{n_s^2}{s} + \frac{n_t^2}{t} + \frac{n_u^2}{u} = \mathcal{M}_4, \quad (18)$$

exactly (15). So the toy dual numerators satisfy the four-point KLT relation on the nose, provided the relative trace sign is kept. The companion code evaluates both sides of (18) over the rationals on a sample of kinematic points and confirms the equality, alongside the color and kinematic Jacobi identities, generalized gauge invariance, crossing symmetry of (15), and the rank-one degeneracy of the biadjoint matrix.

Example 7.2 (Four-point KLT kernel as inverse biadjoint). The biadjoint doubly-partial amplitudes at $n = 4$, on the ordering basis $\{(1234), (1243), (1324)\}$ with planar channel pairs $\{s, t\}, \{s, u\}, \{t, u\}$, form the symmetric matrix

$$m = \begin{pmatrix} \frac{1}{s} + \frac{1}{t} & -\frac{1}{s} & -\frac{1}{t} \\ -\frac{1}{s} & \frac{1}{s} + \frac{1}{u} & -\frac{1}{u} \\ -\frac{1}{t} & -\frac{1}{u} & \frac{1}{t} + \frac{1}{u} \end{pmatrix}.$$

Each row sums to zero, and using $s + t + u = 0$ every 2×2 minor vanishes, so m has rank $(n - 3)! = 1$. The KLT kernel is the inverse of m on this one-dimensional image, and the resulting number depends on which left/right orderings one pairs. Pairing the same ordering (1234) on both sides uses the diagonal entry $m(1234|1234) = 1/s + 1/t = -u/(st)$, whose inverse is $-st/u$, giving the equivalent form $\mathcal{M}_4 = -(st/u) A(1234) \tilde{A}(1234)$. Pairing (1234) on the left with (1243) on the right uses the off-diagonal entry $m(1234|1243) = -1/s$, whose inverse is $-s$, giving the canonical form (16). Both pairings reproduce (15); the companion code checks the numeric equality for the $-s$ pairing and verifies the rank-one degeneracy of m (all nine 2×2 minors vanish), which is exactly the statement that m^{-1} is well-defined on the physical image.

7.6 Five points, sketched

At $n = 5$ there are $(2 \cdot 5 - 5)!! = 15$ cubic graphs. Their propagators are the two-particle Mandelstam invariants s_{ij} , of which there are $\binom{5}{2} = 10$; momentum conservation and masslessness ($\sum_{j \neq i} s_{ij} = 0$ for each i) leave five independent invariants. The Jacobi triples now form an overconstrained-looking but consistent linear system; the minimal basis of partial amplitudes has dimension $(n - 3)! = 2$, so two partial amplitudes—say $A(1, 2, 3, 4, 5)$ and $A(1, 3, 2, 4, 5)$ —determine all others through the fundamental BCJ relations. The double copy (6) becomes a 2×2 KLT bilinear,

$$\mathcal{M}_5 = \sum_{\sigma, \tau \in \{(23), (32)\}} A(\sigma) \mathcal{S}[\sigma|\tau] \tilde{A}(\tau),$$

with the 2×2 kernel $\mathcal{S} = m^{-1}$ the inverse of the 2×2 biadjoint matrix on the same basis. The structure is identical to four points; only the linear algebra grows. We treat five points as a sketch and do not carry the explicit fifteen numerators; the four-point case already exhibits color Jacobi, its kinematic mirror, the double-copy assembly, generalized gauge invariance, and the inverse-biadjoint kernel, which are the whole content of the construction.

8 The classical double copy

The double copy is not confined to amplitudes. A parallel statement relates exact classical solutions of general relativity to exact classical solutions of gauge theory. We summarize the two established forms, keeping their epistemic labels honest.

8.1 The Kerr–Schild double copy

Definition 8.1 (Kerr–Schild metric). A metric is of *Kerr–Schild* form if

$$g_{\mu\nu} = \eta_{\mu\nu} + \phi k_\mu k_\nu,$$

with ϕ a scalar profile and k_μ a vector that is null and geodesic with respect to both η and g : $\eta^{\mu\nu} k_\mu k_\nu = 0$ and $k^\nu \partial_\nu k^\mu = 0$.

The Kerr–Schild ansatz linearizes the Einstein tensor exactly: for such g , the Einstein equations become linear in ϕ . Monteiro, O’Connell and White observed that the same data (ϕ, k) define a gauge field.

Theorem 8.2 (Kerr–Schild double copy). *Let $g_{\mu\nu} = \eta_{\mu\nu} + \phi k_\mu k_\nu$ be a stationary Kerr–Schild solution of the Einstein equations, vacuum away from a localized source, with k null and geodesic. Then the single-copy field $A_\mu^a = c^a \phi k_\mu$, for an arbitrary constant color vector c^a , solves the flat-space Maxwell (indeed abelian Yang–Mills) equations with a source that is the single copy of the gravitational source, and the zeroth-copy scalar ϕ solves the flat-space wave equation with the corresponding scalar source. In the vacuum region both copies solve the source-free equations. [S/H]*

Proof sketch. For Kerr–Schild g the mixed-index Ricci tensor $R^\mu{}_\nu$ is linear in ϕ and its vanishing (or its equality to a source) reduces to a linear equation for $\phi k_\mu k_\nu$. Contracting one index with k and using nullity and the geodesic property collapses the tensor equation to the Maxwell equation for $A_\mu = \phi k_\mu$; contracting both indices collapses it to the scalar wave equation for ϕ . The abelian and non-abelian cases coincide because $A_\mu^a = c^a \phi k_\mu$ has vanishing commutator self-interaction (the color is a constant times a single spacetime vector). Schwarzschild ($\phi \propto 1/r$, k the radial null vector) single-copies to the Coulomb field; Kerr single-copies to the $\sqrt{\text{Kerr}}$ gauge field of a rotating charge disk. $\square$

8.2 The Weyl double copy

The Kerr–Schild double copy is a statement about metrics; the Weyl double copy is a cleaner statement about curvature, in spinor form.

Theorem 8.3 (Weyl double copy, type D). *For a vacuum type D spacetime, the totally symmetric Weyl spinor C_{ABCD} factorizes as*

$$C_{ABCD} = \frac{1}{S} f_{(AB} f_{CD)}, \quad (19)$$

where f_{AB} is a Maxwell field-strength spinor solving the source-free Maxwell equation on the background and S is a scalar (the zeroth copy) solving the scalar wave equation, both on the same spacetime. [S/H]

Later work extended (19) to radiative and more general algebraically special fields (Godazgar, Godazgar, Monteiro, Peinador Veiga and Pope). We flag the limitation: the Weyl double copy in the clean form (19) is established for algebraically special (principally type D) vacua; a background-independent statement for arbitrary curved spacetimes is not available, and the general-purpose status of the classical double copy is therefore [H], not [S].

8.3 Reconstruction reading

The classical double copy makes the reconstruction thesis geometric. A black hole is not, in this reading, a primitive gravitational object. Its curvature is the symmetrized square of a gauge-field-strength spinor divided by a scalar, and its metric is a square of the same null data that builds a Coulomb or $\sqrt{\text{Kerr}}$ gauge field. Spacetime curvature is reconstructed from gauge data exactly as the S-matrix is, using the same squaring invariant.

9 Relation to neighbouring regimes

We locate the double copy among the other QG-IV and QG-adjacent regimes, tracking epistemic status under the worst-component-wins calculus of Section 2.2.

9.1 Positive geometry (companion QG-IV paper)

The seam is Theorem 6.2: the KLT kernel is the inverse of the biadjoint amplitude, which the companion paper produces as the canonical form of the ABHY associahedron. The composite status is

$$\text{worst}([S/H]_{\text{double copy}} \circ [S/H]_{\text{associahedron}}) = [S/H].$$

The two constructions reinforce rather than degrade one another: both halves of QG-IV are S/H , and their shared invariant (the kernel) is S as mathematics.

9.2 Motivic coaction and open problem 1

Attaching the double-copy amplitude to its transcendental (stringy or dimensionally regulated) completion brings in the motivic coaction of the amplitude’s period, and the central conjecture “residue tree = coaction tree” (open problem 1). The unitarity cuts of a double-copy loop integrand are the primary tool for building and checking BCJ-dual loop numerators, so the double copy and the coaction interact directly. But the coaction side is [H] (it assumes a mixed-Tate regime; elliptic and non-Tate sectors need enlarged coalgebras), and the tree-identification is [H] in verified low-weight cases and [P] in general. Hence

$$\text{worst}([S/H]_{\text{double copy}} \circ [H]_{\text{coaction}} \circ [H/P]_{\text{residue=coaction}}) = [H/P].$$

The double copy cannot rescue the coaction conjecture; the weakest link governs.

9.3 Cut/regulator complementarity and open problem 2

Open problem 2 asks whether discontinuity/cut maps are adjoint to regulator/period maps in a suitable category. Because unitarity cuts drive the loop-level double copy, a precise cut/regulator adjunction would immediately sharpen the loop-level existence question of Section 10. We record this as a target, not a result; its status is [H].

9.4 Celestial holography

Mellin-transforming a double-copy S-matrix to the celestial sphere produces celestial correlators; the “celestial double copy” is an active subject. The celestial dictionary is [H] (its soft-sector and asymptotic-symmetry assumptions are not settled), so the composite with the double copy is [H].

9.5 Summary table

Composition	Shared invariant	Composite status
Double copy $\circ$ associahedron	KLT kernel $= m^{-1}$	[S/H]
Double copy $\circ$ coaction $\circ$ residue=coaction	residue/cut tree	[H/P]
Double copy $\circ$ cut/regulator adjunction	unitarity cut	[H]
Double copy $\circ$ celestial basis	Mellin S-matrix	[H]

10 Honest limitations and open problems

We state the obstructions plainly; none is hidden by the slogan.

No general all-loop existence proof. Color–kinematics dual numerators are guaranteed at tree level (Theorem 4.3). At loop level their existence is established case by case—famously through four loops in the four-point amplitude of $\mathcal{N} = 8$ supergravity and to high order in $\mathcal{N} = 4$ super-Yang–Mills—but there is no general theorem that a BCJ-dual representation exists at arbitrary loop order and multiplicity. The kinematic-algebra program (Brandhuber–Chen–Johansson–Travaglini–Wen and successors) gives closed forms in specific sectors, not in general. Status of the loop-level double copy as physics: [H].

Locality versus duality at loop level. Even when dual loop numerators exist, they may be non-local (contain spurious poles) if one insists on manifest duality, or local at the cost of manifest duality. Reconciling the two is an open technical problem and is one reason the loop-level statements resist a clean existence theorem.

The square is not pure gravity. Yang–Mills squared is $\mathcal{N} = 0$ supergravity—graviton plus dilaton plus two-form, not pure Einstein gravity. Extracting pure gravity requires ghost

matter or projection. Any statement “gravity = gauge²” that omits the extra states is imprecise; we have kept them explicit.

The classical double copy is not background independent. The Kerr–Schild and Weyl double copies are cleanest for algebraically special (Kerr–Schild, type D) vacua. A general statement for arbitrary curved backgrounds is not established; attempts to extend beyond algebraically special solutions meet genuine obstructions. Status: [H] as a general claim.

Off-shell / Lagrangian-level duality. A local action manifesting color–kinematics duality to all orders in the fields is not known; the duality is manifest on-shell and, in specific constructions, up to a finite order off-shell. This bears on any attempt to make the double copy a statement about actions rather than amplitudes.

Open problems from the roadmap. Item 1 (residue tree = coaction tree) and item 2 (cut/regulator complementarity) touch this regime directly, as discussed in Section 9. Neither is resolved here; both are flagged [H] or [H/P].

11 The companion Haskell development

The paper is accompanied by a small Haskell package under `src/double-copy/` that makes the four-point construction executable and its algebraic claims checkable over exact rational arithmetic. It has no dependencies outside `base` and compiles with GHC.

Modules. `Core.hs` defines the three four-point channels, Mandelstam data over $\mathbb{Q}$, the $\mathfrak{su}(2)$ structure constants ε^{abc} , the color factors c_s, c_t, c_u of Theorem 3.3, the toy color-dual numerators (14), the double-copy assembly (15), the KLT bilinear (16), and the biadjoint doubly-partial amplitudes of Theorem 6.1. `Properties.hs` states the checkable properties. `Main.hs` prints the demonstrations and runs the checks, exiting nonzero on any failure.

Checkable properties. The package verifies, over exact rationals and for a deterministic sample of kinematic points:

1. **Color Jacobi.** $c_s + c_t + c_u = 0$ as an identity of $\mathfrak{su}(2)$ tensors, checked for all 3^4 external-index assignments (Theorem 3.2, Theorem 3.3).
2. **Mandelstam constraint.** $s + t + u = 0$ for the generated kinematics ((10)).
3. **Kinematic Jacobi.** $n_s + n_t + n_u = 0$ for the toy dual numerators (14) (Theorem 4.1).
4. **Generalized gauge invariance.** The partial amplitude (12) and the full amplitude are invariant under the shift $n_g \mapsto n_g + D_g \alpha$ (Theorem 7.1), the amplitude-level shadow of color Jacobi.

5. **Crossing symmetry of the double copy.** The double-copy amplitude (15) is invariant under permutations of (s, t, u) with the induced numerator permutation.
6. **KLT match.** The four-point KLT bilinear $-s A(1, 2, 3, 4) \tilde{A}(1, 2, 4, 3)$, built from the partials (12) with their trace signs, equals the direct double-copy assembly (15) on the sampled kinematics ((18)).
7. **Biadjoint rank one.** On the ordering basis $\{(1234), (1243), (1324)\}$ all nine 2×2 minors of the biadjoint matrix m vanish (using $s + t + u = 0$) while m is nonzero, so m has rank $(n - 3)! = 1$ and the KLT kernel m^{-1} is well-defined on the physical image (Theorem 7.2).
8. **Reduced assembly.** The double copy (15) agrees with its u -eliminated form (with $u = -(s + t)$) on the sampled kinematics.

The kinematic Jacobi check (item 3) is the property required by the project’s build gate: it is exact, has no free parameters once the dual set is fixed, and fails loudly if the numerator algebra is perturbed.

12 Conclusion

The double copy realizes the project’s reconstruction thesis with unusual directness. The invariant is a gauge-theory kinematic numerator constrained to obey the Lie-algebra Jacobi identities of color; gravity, both its S-matrix and—through the classical double copy—its curvature, is the algebraic square of that invariant. Nothing gravitational is posited beyond the choice of pairing. We have given the cubic-graph organization, the color Jacobi relations and their kinematic mirror, the double-copy substitution, the KLT relations, and the identification of the KLT kernel with the inverse biadjoint amplitude that stitches this paper to its positive-geometry companion. We worked four points in full and sketched five, and we treated the classical Kerr–Schild and Weyl double copies. Throughout we kept the epistemic labels honest: tree-level amplitude mathematics is S ; its reading as a physical gravity S-matrix is S/H ; the loop-level and general classical statements are H ; and the conjectural bridge to motivic coactions is H/P . The construction is not a theory of quantum gravity. It is a precise demonstration that, in the perturbative regime, gravity is not an independent input but a reconstruction—a square—of gauge-theoretic representation data, with the KLT kernel as the sharp invariant shared across the amplitude module of the library.

References

- [1] Z. Bern, J. J. M. Carrasco, and H. Johansson, “New Relations for Gauge-Theory Amplitudes,” *Phys. Rev. D* **78** (2008) 085011; arXiv:0805.3993 [hep-ph].
- [2] Z. Bern, J. J. M. Carrasco, and H. Johansson, “Perturbative Quantum Gravity as a Double Copy of Gauge Theory,” *Phys. Rev. Lett.* **105** (2010) 061602; arXiv:1004.0476 [hep-th].

- [3] H. Kawai, D. C. Lewellen, and S. H. H. Tye, “A Relation Between Tree Amplitudes of Closed and Open Strings,” *Nucl. Phys. B* **269** (1986) 1–23.
- [4] N. E. J. Bjerrum-Bohr, P. H. Damgaard, and P. Vanhove, “Minimal Basis for Gauge Theory Amplitudes,” *Phys. Rev. Lett.* **103** (2009) 161602; arXiv:0907.1425 [hep-th].
- [5] N. E. J. Bjerrum-Bohr, P. H. Damgaard, T. Søndergaard, and P. Vanhove, “The Momentum Kernel of Gauge and Gravity Theories,” *JHEP* **01** (2011) 001; arXiv:1010.3933 [hep-th].
- [6] F. Cachazo, S. He, and E. Y. Yuan, “Scattering of Massless Particles in Arbitrary Dimensions,” *Phys. Rev. Lett.* **113** (2014) 171601; arXiv:1307.2199 [hep-th].
- [7] F. Cachazo, S. He, and E. Y. Yuan, “Scattering equations and Kawai-Lewellen-Tye orthogonality,” *Phys. Rev. D* **90** (2014) 065001; arXiv:1306.6575 [hep-th].
- [8] S. Mizera, “Combinatorics and topology of Kawai-Lewellen-Tye relations,” *JHEP* **08** (2017) 097; arXiv:1706.08527 [hep-th].
- [9] S. Mizera, “Inverse of the String Theory KLT Kernel,” *JHEP* **06** (2017) 084; arXiv:1610.04230 [hep-th].
- [10] Z. Bern, J. J. M. Carrasco, M. Chiodaroli, H. Johansson, and R. Roiban, “The SAGEX Review on Scattering Amplitudes, Chapter 2: An Invitation to Color-Kinematics Duality and the Double Copy,” *J. Phys. A* **55** (2022) 443003; arXiv:2203.13013 [hep-th].
- [11] Z. Bern, J. J. M. Carrasco, M. Chiodaroli, H. Johansson, and R. Roiban, “The Duality Between Color and Kinematics and its Applications,” arXiv:1909.01358 [hep-th].
- [12] R. Monteiro, D. O’Connell, and C. D. White, “Black holes and the double copy,” *JHEP* **12** (2014) 056; arXiv:1410.0239 [hep-th].
- [13] A. Luna, R. Monteiro, I. Nicholson, and D. O’Connell, “Type D Spacetimes and the Weyl Double Copy,” *Class. Quantum Grav.* **36** (2019) 065003; arXiv:1810.08183 [hep-th].
- [14] M. Godazgar, K. Godazgar, R. Monteiro, D. Peinador Veiga, and C. N. Pope, “Weyl Double Copy for Gravitational Waves,” *Phys. Rev. Lett.* **126** (2021) 101103; arXiv:2010.02925 [hep-th].
- [15] A. Brandhuber, G. Chen, H. Johansson, G. Travaglini, and C. Wen, “Kinematic Hopf Algebra for Bern-Carrasco-Johansson Numerators in Heavy-Mass Effective Field Theory and Yang-Mills Theory,” *Phys. Rev. Lett.* **128** (2022) 121601; arXiv:2111.15649 [hep-th].
- [16] C. Cheung, C.-H. Shen, and C. Wen, “Unifying Relations for Scattering Amplitudes,” *JHEP* **02** (2018) 095; arXiv:1705.03025 [hep-th].
- [17] C. D. White, “Snowmass White Paper: the Double Copy and its Applications,” arXiv:2204.06547 [hep-th].

- [18] Z. Bern, J. J. M. Carrasco, L. J. Dixon, H. Johansson, and R. Roiban, “Manifest Ultraviolet Behavior for the Three-Loop Four-Point Amplitude of $\mathcal{N} = 8$ Supergravity,” *Phys. Rev. D* **78** (2008) 105019; arXiv:0808.4112 [hep-th].
- [19] Z. Bern, L. J. Dixon, D. C. Dunbar, M. Perelstein, and J. S. Rozowsky, “Multi-Leg One-Loop Gravity Amplitudes from Gauge Theory,” *Nucl. Phys. B* **546** (1999) 423; hep-th/9811140.
- [20] R. Monteiro and D. O’Connell, “The Kinematic Algebra From the Self-Dual Sector,” *JHEP* **07** (2011) 007; arXiv:1105.2565 [hep-th].
- [21] N. Arkani-Hamed, Y. Bai, S. He, and G. Yan, “Scattering Forms and the Positive Geometry of Kinematics, Color and the Worldsheet,” *JHEP* **05** (2018) 096; arXiv:1711.09102 [hep-th].