

Durham E-Theses

Precision QCD corrections to gluon-initiated diphoton-plus-jet production at the LHC

MOODIE, RYAN, IAIN

How to cite:

MOODIE, RYAN, IAIN (2022) *Precision QCD corrections to gluon-initiated diphoton-plus-jet production at the LHC*, Durham theses, Durham University. Available at Durham E-Theses Online: <http://etheses.dur.ac.uk/14653/>

Use policy



This work is licensed under a [Creative Commons Attribution 3.0 \(CC BY\)](https://creativecommons.org/licenses/by/3.0/)

Precision QCD corrections to gluon-initiated diphoton-plus-jet production at the LHC

Ryan Iain Moodie

A Thesis presented for the degree of
Doctor of Philosophy



Institute for Particle Physics Phenomenology
Department of Physics
Durham University
United Kingdom

April 2022

Precision QCD corrections to gluon-initiated diphoton-plus-jet production at the LHC

Ryan Iain Moodie

Submitted for the degree of Doctor of Philosophy

April 2022

Abstract: In this thesis, we present recent advances at the precision frontier of higher-order quantum chromodynamics (QCD) calculations. We consider massless two-loop five-point amplitudes, with a particular focus on diphoton-plus-jet production through gluon fusion. We build a library of infrared functions up to at most next-to-next-to-leading order (NNLO) in QCD, which can be used to validate amplitudes and construct counterterms in subtraction schemes at NNLO. We review progress in the novel use of machine learning technology to optimise the evaluation of amplitudes in hadron collider simulations. We present the full-colour virtual QCD corrections to diphoton-plus-jet production through gluon fusion, discussing the new techniques developed to calculate these non-planar two-loop amplitudes. We use these amplitudes to compute the next-to-leading order QCD corrections to the differential cross sections of diphoton-plus-jet production through gluon fusion at the Large Hadron Collider. We also present the leading-colour double-virtual corrections to hadronic trijet production. All derived amplitudes are made available in a public implementation that is ready for further phenomenological application.

Contents

Abstract	ii
List of Abbreviations	vi
1 Introduction	1
1.1 The Standard Model	3
1.1.1 Quantum electrodynamics	3
1.1.2 Quantum chromodynamics	5
1.2 Observables	8
1.2.1 Factorisation theorem	8
1.2.2 Cross sections	9
1.2.3 Beyond fixed order	11
1.3 Amplitudes	12
1.3.1 Properties	13
1.3.2 Infinities	14
1.4 Colour	15
1.4.1 Colour decomposition	15
1.4.2 Colour matrices	19
1.5 Kinematics	22
1.5.1 Degrees of freedom	23
1.5.2 Representations	25
1.5.3 Weyl spinors	25
1.5.4 Spinor-helicity formalism	27
1.5.5 Helicity amplitudes	31
1.5.6 Recursion	32
1.5.7 Momentum twistor variables	33
1.6 On-shell methods and integral reduction	35
1.6.1 Loops	35
1.6.2 Unitarity cuts	36
1.6.3 Generalised unitarity	38
1.6.4 Integrand reduction	41

1.6.5	Integration-by-parts identities	42
1.7	Finite fields	44
1.7.1	Definition	44
1.7.2	Reconstruction	45
1.7.3	Recovering rationals	47
1.8	Phenomenology	48
1.8.1	Fixed-order perturbation theory	49
1.8.2	Estimating uncertainty	51
1.8.3	Diphoton production	54
1.9	Structure of this thesis	55
2	Infrared behaviour	56
2.1	Infrared limits	56
2.1.1	Cuts	57
2.1.2	Poles	58
2.2	Infrared factorisation	58
2.2.1	Collinear limits	59
2.2.2	Soft limits	61
2.2.3	Beyond NLO	64
2.2.4	Utility	65
2.3	Infrared subtraction	65
2.4	Implementation	67
2.4.1	Validation	68
2.4.2	Example code	71
3	Matrix element neural networks	78
3.1	Background	78
3.2	Amplitudes	79
3.3	Phase space partitioning	81
3.4	Model architecture	82
3.5	Pipeline	83
3.6	Results	85
3.7	Summary	88
4	Virtual QCD corrections to the $gg \rightarrow \gamma\gamma g$ amplitude	90
4.1	Background	91
4.2	Kinematics and amplitude conventions	92
4.3	Computational setup and amplitude reduction	95
4.4	Momentum twistor parametrisation	98
4.5	Analytic reconstruction over finite fields	100

4.5.1	Linear relations among the rational coefficients	100
4.5.2	Matching factors on univariate slices	101
4.5.3	Univariate partial fraction decomposition over finite fields	102
4.5.4	Summary and impact of the reconstruction strategy	106
4.6	Compact analytic expressions for the all-plus configuration	108
4.7	Implementation and performance	110
4.8	Summary	115
5	NLO QCD corrections to $gg \rightarrow \gamma\gamma g$ distributions	117
5.1	Background	117
5.2	Computational setup	118
5.3	Results	120
5.4	Computational cost	128
5.5	Summary	129
6	Leading-colour double-virtual QCD corrections to $pp \rightarrow 3j$	130
6.1	Stability and timing	130
6.2	Infrared performance	132
7	Conclusion	134
	Bibliography	136
	Index	166

List of Abbreviations

API	application programming interface	83
ATLAS	<i>A Toroidal LHC Apparatus</i>	48
BCFW	Britto-Cachazo-Feng-Witten	33
BLHA	<i>Binoth Les Houches Accord</i>	85
CMS	<i>Compact Muon Solenoid</i>	48
CPU	central processing unit	81
CS	Catani-Seymour	59
DGLAP	Dokshitzer-Gribov-Lipatov-Altarelli-Parisi	8
DR	dimensional regularisation	14
EM	electromagnetic	3
ESSB	electroweak spontaneous symmetry breaking	3
f32	32-bit floating-point number (single precision)	83
f64	64-bit floating-point number (double precision)	67
f128	128-bit floating-point number (quadruple precision)	67
f256	256-bit floating-point number (octuple precision)	67
FC	full-colour	54
FF	finite field	35
FKS	Frixione-Kunszt-Signer	66
FR	finite remainder	66
FS	full-spin	61
IBP	integration-by-parts	43
IR	infrared	11
ISP	irreducible scalar product	41
KK	Kleiss-Kuijf	18

KLN	Kinoshita-Lee-Nauenberg	58
LC	leading-colour	19
LHC	Large Hadron Collider	9
LI	linearly independent	18
LO	leading order	13
LS	leading-spin	61
MC	Monte Carlo	9
ME	matrix element	9
$\overline{\text{MHV}}$	anti-maximally-helicity-violating	31
MHV	maximally-helicity-violating	31
MI	master integral	42
ML	machine learning	78
$\overline{\text{MS}}$	modified minimal subtraction	15
MSE	mean squared error	83
MTV	momentum twistor variable	34
N³LO	(next-to) ³ leading order	50
N⁴LO	(next-to) ⁴ leading order	54
N^kLL	(next-to) ^k leading log	11
N^kLO	(next-to) ^k leading order	51
N^kLP	(next-to) ^k leading power	11
NLO	next-to-leading order	32
NMHV	next-to-maximally-helicity-violating	31
NNLO	next-to-next-to-leading order	32
NN	neural network	55
OPP	Ossola-Pittau-Papadopoulos	42
PDF	parton distribution function	8
PFD	partial fraction decomposition	58
QCD	quantum chromodynamics	1
QED	quantum electrodynamics	3
QFT	quantum field theory	3
RMSE	root mean squared error	86

R³	triple-real	50
RRV	real-real-virtual	50
RR	double-real	50
RVV	real-virtual-virtual	50
RV	real-virtual	50
SCET	Soft-Collinear Effective Theory	11
Sherpa	<i>Simulation of High Energy Reactions of Particles</i>	78
SLC	subleading-colour	93
SM	Standard Model	1
tHV	't Hooft-Veltman	15
UHV	ultra-helicity-violating	32
UV	ultraviolet	14
V³	triple-virtual	50
VV	double-virtual	32

Declaration

The work in this thesis is based on research carried out in the Institute for Particle Physics Phenomenology at Durham University. No part of this thesis has been submitted elsewhere for any degree or qualification. This thesis is based on joint research: Chapter 1 is a review of the literature; Chapter 2 discusses contributions to the C++ amplitude library `NJet3` [1]; Chapter 3 is based on Ref. [2]; Chapter 4 on Ref. [3]; Chapter 5 on Ref. [4]; and Chapter 6 on Ref. [5].

Copyright © 2022 Ryan Iain Moodie.

The copyright of this thesis rests with the author. No quotation from it should be published without the author's prior written consent and information derived from it should be acknowledged.

Acknowledgements

Foremost thanks to my supervisor, Simon Badger, for his mentorship over the past few years. Thanks to all those I've had the pleasure of collaborating with and to everyone in the IPPP community that I've enjoyed being a part of so much. For proofreading, helpful comments, and discussions on the draft, a huge thanks to Oscar Braun-White, Lucy Budge, Lois Flower, Hitham Hassan, Sebastian Jaskiewicz, Stephen Jones, Matteo Marcoli, Francesco Sarandrea, Henry Truong, Yannick Ulrich, Mia West, and Simone Zoia.

The work presented in this thesis was funded by UKRI-STFC grant numbers ST/S505365/1 and ST/P001246/1.

Chapter 1

Introduction

The Standard Model (SM) is one of the most successful scientific theories of history, including the most precise agreements of prediction and experiment ever achieved [6, 7]. However, it fails to describe some observed phenomena of the universe and contains empirical parameters [8]. The subject of this thesis is precision quantum chromodynamics (QCD) phenomenology, which is an indirect search for physics beyond the SM by comparing high-precision SM predictions to measurements at hadron colliders [9]. New physics would appear as small deviations from SM expectations [10].

In this chapter, we introduce the basic toolkit required to calculate QCD corrections to amplitudes and use them to construct high-precision predictions for observables at colliders. We begin in Section 1.1 with a brief review of the relevant sectors of the SM. We discuss how amplitudes can be used to construct collider observables in Section 1.2, touching on some details of amplitudes in Section 1.3. In Section 1.4, we look at how we can simplify the computation of QCD amplitudes by factorising into colour and kinematic contributions, treating the former with colour decomposition. In Section 1.5, we cover various representations of the kinematics and some methods for computing partial amplitudes. In Section 1.6, we discuss how on-shell methods can be used to calculate loop-level amplitudes. In Section 1.7, we show how finite field arithmetic can aid the computation of amplitudes. We motivate the phenomenology of diphoton production in Section 1.8, before summarising the contents of this thesis in Section 1.9.

For further reading on these topics, see the textbooks [11–21], lectures [22–24], reviews [25–28], and theses [29–31].




Figure 1.1: Table of the particle content of the SM. Outlines of the corresponding colour contain particles carrying strong (red), electromagnetic (blue), or weak (green) gauge charge. The electric charge (blue) and colour representation (red) are shown in the upper right corner of appropriate particle boxes. For massive particles, there is a yellow circle with area proportional to the mass on a square-root scale (data from Ref. [8]).

1.1 The Standard Model

The SM is our current best description of the fundamental structure of the universe. It is a quantum field theory (QFT) with a direct product gauge group,

$$SU_C(3) \times SU_L(2) \times U_Y(1) \rightarrow SU_C(3) \times U_{EM}(1). \quad (1.1.1)$$

The first Lie group is the strong interaction, with subscript C for colour, which is described by QCD. The next two comprise the electroweak sector, with subscript L for left-handed chirality and Y for hypercharge. The transition in Eq. (1.1.1) is electroweak spontaneous symmetry breaking (ESSB) through the Higgs mechanism [32, 33], which gives rise to the weak and electromagnetic (EM) interactions we observe. The resultant EM group is described by quantum electrodynamics (QED). The SM does not describe gravity.

The forces are mediated by vector bosons (spin-1), also called gauge bosons: the massless gluon g for the strong force, the massive W^\pm and Z bosons for the weak force, and the massless photon γ for the EM force. The matter particles of the SM are fermions (spin-1/2). This includes the quarks—the up u , down d , charm c , strange s , top t , and bottom b —which are massive and carry colour, weak charge, and electric charge. The other fermions are the leptons, including: charged leptons—the electron e , muon μ , and tau τ —which are massive and carry weak and electric charge; and the neutrinos—the electron neutrino ν_e , muon neutrino ν_μ , and tau neutrino ν_τ —which are massless¹ and carry only weak charge. The different types of quarks and leptons are referred to as flavours. Each fermion comes with an antiparticle, which we denote with an overline. There is also the Higgs boson H , which is a scalar boson (spin-0). It generates the masses of the W^\pm and Z bosons, the charged leptons, and the quarks through ESSB. The particles of the SM are tabulated in Fig. 1.1.

As QED and QCD are the relevant sectors for this thesis, we will further discuss them in the following sections. The complete SM Lagrangian, along with all Feynman rules, is presented in Ref. [34].

1.1.1 Quantum electrodynamics

QED is a QFT describing the EM interaction. It is an abelian gauge theory with gauge group $U(1)$. The gauge charge is electric charge. The photon field A_μ couples to (anti)fermion fields $(\overline{\psi}_f) \psi_f$ with mass m_f , which includes the quarks and charged

¹While neutrinos are massless in the SM, this is a shortcoming of the model as the experimental evidence of neutrino oscillations indicates that they have a small yet non-zero mass.

leptons. With the imaginary unit,

$$i^2 = -1, \quad (1.1.2)$$

and the Dirac adjoint and Dirac slash,

$$\begin{aligned} \bar{\psi} &= \psi^\dagger \gamma^0, \\ \not{D} &= D_\mu \gamma^\mu, \end{aligned} \quad (1.1.3)$$

where γ^μ with $\mu \in \{0, 1, 2, 3\}$ in four dimensions are the gamma matrices, the classical part of the QED Lagrangian is

$$\mathcal{L}_{\text{QED}}^{\text{classical}} = -\frac{1}{4} F_{\mu\nu} F^{\mu\nu} + \sum_f \bar{\psi}_f (i\not{D} - m_f) \psi_f, \quad (1.1.4)$$

with the field strength tensor,

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu, \quad (1.1.5)$$

and the covariant derivative,

$$D_\mu = \partial_\mu - iQ_f e A_\mu. \quad (1.1.6)$$

The EM coupling $Q_f e$ is equal to the electric charge of the fermion f , with

$$Q_f = \begin{cases} -1 & \text{charged lepton} & (e, \mu, \tau), \\ \frac{2}{3} & \text{up-type quark} & (u, c, t), \\ -\frac{1}{3} & \text{down-type quark} & (d, s, b). \end{cases} \quad (1.1.7)$$

The coupling is often expressed in terms of the dimensionless quantity,

$$\alpha = \frac{e^2}{4\pi}, \quad (1.1.8)$$

called the fine structure constant. The coupling depends on the energy scale, as discussed in Section 1.1.2.

The quantisation of the Lagrangian involves adding a gauge fixing term. This allows a propagator to be defined for the photon, which depends on the gauge choice, such as R_ξ gauge (Chapter 62 of Ref. [15]). Because the gauge degrees of freedom are unphysical, any physical predictions of the theory are independent of the choice of gauge. They are also unchanged by gauge transformations, which is a property we call gauge invariance. We discuss gauges in more detail for the non-abelian case of QCD.

1.1.2 Quantum chromodynamics

QCD is a QFT describing the strong interaction. It is a non-abelian gauge theory with gauge group $SU(3)$. The gauge charge is called colour charge. The gluon fields A_μ^a lie in the adjoint representation with $a \in \{1, \dots, 8\}$, and are coupled to (anti)quark fields $(\overline{\psi}_q^i) \psi_q^i$ lying in the (anti)fundamental representation with $i \in \{1, 2, 3\}$. The classical part of the QCD Lagrangian is

$$\begin{aligned} \mathcal{L}_{\text{QCD}}^{\text{classical}} &= -\frac{1}{4} F_{\mu\nu}^a F^{a,\mu\nu} + \sum_q \overline{\psi}_q^i (\mathbf{i} \not{D}_{ij} - \delta_{ij} m_q) \psi_q^j, \\ F_{\mu\nu}^a &= \partial_\mu A_\nu^a - \partial_\nu A_\mu^a + g_s f^{abc} A_\mu^b A_\nu^c, \\ (D_\mu)_{ij} &= \delta_{ij} \partial_\mu - \mathbf{i} g_s t_{ij}^a A_\mu^a, \end{aligned} \tag{1.1.9}$$

where t_{ij}^a are the $SU(3)$ generators in a fundamental representation. There exist various possible matrix representations; one is through the proportionality to the Gell-Mann matrices [35] λ_{ij}^a ,

$$t_{ij}^a = \frac{1}{2} \lambda_{ij}^a. \tag{1.1.10}$$

The structure constants f^{abc} are defined by the commutator of the fundamental generators,

$$[t^a, t^b] = \mathbf{i} f^{abc} t^c. \tag{1.1.11}$$

The strong coupling g_s is often expressed in terms of the dimensionless parameter,

$$\alpha_s = \frac{g_s^2}{4\pi}. \tag{1.1.12}$$

Again, a gauge fixing term is introduced as in the abelian case of QED. R_ξ gauge adds a covariant gauge fixing term,

$$\mathcal{L}_{\text{QCD}}^{R_\xi} = -\frac{1}{2\xi} (\partial^\mu A_\mu^a)^2, \tag{1.1.13}$$

which leads to the gluon propagator,

$$G_{\mu\nu}^{ab}(p) = \frac{\mathbf{i} \delta^{ab}}{p^2 + \mathbf{i}\varepsilon} \left(-\eta_{\mu\nu} + (1 - \xi) \frac{p_\mu p_\nu}{p^2 + \mathbf{i}\varepsilon} \right). \tag{1.1.14}$$

This encompasses Feynman gauge for $\xi = 1$ and Landau gauge for $\xi \rightarrow 0$.

With covariant gauges, the non-abelian theory can also require the introduction of ghost fields to cancel unphysical modes through a procedure such as the Faddeev-Popov method [36]. The ghost field c^a is in the adjoint representation of $SU(N_c)$

and is fermionic. The ghost term of the Lagrangian can take the form,

$$\mathcal{L}_{\text{QCD}}^{\text{ghost}} = \left(\partial_\mu \bar{c}^a \right) D_{ab}^\mu c^b, \quad (1.1.15a)$$

$$D_{ab}^\mu = \partial^\mu \delta_{ab} + i(F^c)_{ab}(A^c)^\mu, \quad (1.1.15b)$$

$$(F^a)_{bc} = -if^{abc}, \quad (1.1.15c)$$

where D_{ab}^μ is the covariant derivative in the adjoint representation and $(F^a)_{bc}$ is the $SU(N_c)$ generator in an adjoint representation. These ghosts appear in amplitude loops (Section 1.3).

Axial gauges are an alternative choice which can be ghost-free at the cost of a more complicated gluon propagator. We fix the gauge field with respect to an arbitrary reference vector q^μ ,

$$\mathcal{L}_{\text{QCD}}^{\text{axial}} = -\frac{1}{2\xi} (q^\mu A_\mu^a)^2, \quad (1.1.16)$$

giving the gluon propagator,

$$G_{\mu\nu}^{ab}(p) = \frac{i\delta^{ab}}{p^2 + i\varepsilon} \left(-\eta_{\mu\nu} + \frac{p_\mu q_\nu + p_\nu q_\mu}{p \cdot q} - (q^2 + \xi p^2) \frac{p_\mu p_\nu}{(p \cdot q)^2} \right). \quad (1.1.17)$$

Light-like axial gauge, or light-cone gauge, is given in the case of a null (light-like) reference vector $q^2 = 0$ and $\xi \rightarrow 0$. For $\xi \rightarrow 0$, the ghost fields decouple and can be neglected: the two physical polarisation states of the gluon are explicit.

Couplings depend on the energy scale μ of the physical process [37]. This phenomenon is called the running of the coupling and is described by a Callan-Symanzik β -function [38,39], also called a renormalisation group flow rate. The QCD β -function is known to five loops [40,41] and takes the form

$$\frac{d\alpha_s}{d \ln(\mu^2)} =: \beta(\alpha_s) = -\sum_{n=0}^{\infty} \beta_n \left(\frac{\alpha_s}{4\pi} \right)^{n+2}. \quad (1.1.18)$$

Since it comes with a negative sign, QCD exhibits asymptotic freedom [42,43]: the strong coupling decreases with increasing energy, as shown in Fig. 1.2. The scale separating the strongly- and weakly-coupled QCD phases is $\Lambda_{\text{QCD}} \sim 200$ MeV, with $\alpha_s < 1$ for energy scales $Q \gg \Lambda_{\text{QCD}}$. This means that hard processes—those at high energy or equivalently small length scale—can be treated perturbatively with α_s , Eq. (1.1.12), as the small expansion parameter. Conversely, soft physics must be treated non-perturbatively, as we discuss in Section 1.2.3. Composite QCD states in the bound regime are called hadrons. This includes mesons with an even number of valence quarks such as the pions, and baryons with an odd number of valence quarks such as the proton.




Figure 1.2: Summary of measurements of α_s as a function of the energy scale Q . The degree of QCD perturbation theory (Section 1.8.1) used to extract α_s is indicated in brackets, with +res meaning matched to a resummed calculation (Section 1.2.3). The theory prediction of α_s running is given by the black line and confidence interval. Figure from Ref. [8].

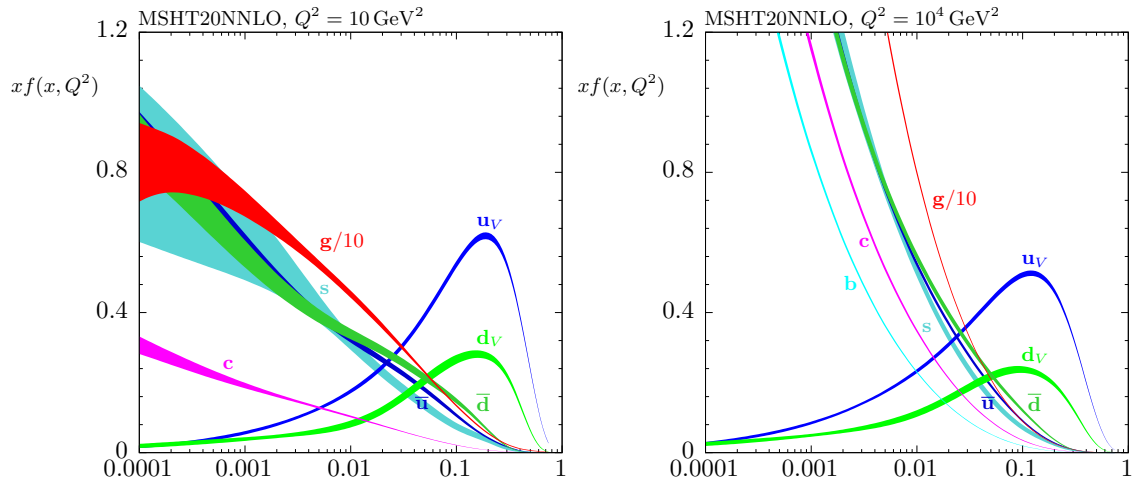


Figure 1.3: MSHT20 NNLO PDFs for the proton at low (left) and high (right) energy scales. The x -axis shows the momentum fraction x . Lines are labelled by their particle type, with subscript V for valence. The gluon PDF is scaled down by a factor of ten. Figure from Ref. [45].

1.2 Observables

In this section, we discuss how to bridge the gap between perturbative QFT calculations, which generally use Feynman diagram technology to compute amplitudes within a theory defined by its Lagrangian, and physical observables, which can be experimentally measured.

1.2.1 Factorisation theorem

The principles of scale factorisation [17,44]—the separation of hard and soft physics—and asymptotic freedom (Section 1.1.2) allow the calculation of observables using QCD at hadron colliders. We introduce the factorisation scale μ_F as the scale of the interface between soft and hard physics.

The soft physics of the composite initial states can be described by a parton distribution function (PDF) [46–48] $f_i(x, \mu_F)$, which is the probability of finding a parton i , meaning a quark or a gluon, with momentum fraction x inside a hadron. PDFs are experimentally determined and evolved between different scales through the Dokshitzer-Gribov-Lipatov-Altarelli-Parisi (DGLAP) equations, just as the running of α_s is described by the QCD β -function (Section 1.1.2). Various PDF sets are available, including Refs. [49–52] and those shown in Fig. 1.3 [45], with recent PDF sets achieving percent-level uncertainties [49]. While the precision is not currently competitive with empirical techniques, there is also the future possibility to obtain the PDFs from first principles through lattice QCD [53].

The hard scattering of the partons is treated in fixed-order perturbation theory

(Section 1.8.1), in which it is given by the squared amplitude (Section 1.3), which we also refer to as the matrix element (ME). To calculate a differential cross section $d\sigma$, the PDFs are convolved with the ME and integrated over the phase space by a Monte Carlo (MC) event generator [54, 55]. There are several available general-purpose MC event generators [56–63]. They additionally simulate the non-perturbative effects discussed in Section 1.2.3.

For a proton-proton collider like the Large Hadron Collider (LHC), the hadronic cross section can be written as

$$d\sigma = \sum_{i,j} \int_0^1 dx_1 dx_2 f_i(x_1, \mu_F) f_j(x_2, \mu_F) d\hat{\sigma}_{ij}(Q, \mu_F, \mu_R) + \mathcal{O}\left(\frac{\Lambda_{\text{QCD}}}{Q}\right). \quad (1.2.1)$$

The partonic, or hard, cross section for the scattering of the partons i and j in the initial states, $d\hat{\sigma}_{ij}$, depends on the partonic energy scale Q , which is given by the partonic centre of mass energy. The renormalisation scale μ_R (Section 1.3) and factorisation scale μ_F appear, although the scale dependence would vanish in an all-orders expression. Further non-perturbative effects can be neglected for $Q \gg \Lambda_{\text{QCD}}$ [64]. The factorisation of partonic and hadronic physics is depicted in Fig. 1.4.

1.2.2 Cross sections

For a two-particle initial state \mathcal{I} and an N -particle final state \mathcal{F} , the hard cross section is given by

$$d\hat{\sigma} = \frac{1}{F} \overline{|\mathcal{A}_{2+N}|^2} d\Phi_N, \quad (1.2.2)$$

where $\overline{|\mathcal{A}_{2+N}|^2}$ is the $2 \rightarrow N$ amplitude squared with the overline denoting that initial spin and colour states are averaged and final spin and colour states are summed, $d\Phi_N$ is the N -particle Lorentz invariant phase space, and F is the flux factor, which is

$$F = 2 s_{12} \quad (1.2.3)$$

for a two-particle initial state ($s_{ij} = (p_i + p_j)^2$, see Section 1.5.1). The treatment of colour derives from the phenomenon of colour confinement, which is that colour is unobservable because only colourless bound hadrons exist at low energy scales $Q < \Lambda_{\text{QCD}}$ (Section 1.1.2). The treatment of spin is on the assumption that the collider uses an unpolarised input and does not measure the polarisation of the product states. The phase space [65] $d\Phi_N$ in d -dimensions has

$$(d-1)N - d \quad (1.2.4)$$




Figure 1.4: Sketch of a hadron collider simulation. A hadron h_1 with momentum P_1 collides with another hadron h_2 with momentum P_2 . A parton i with momentum fraction x_1 is factorised from h_1 through a PDF $f_i(x_1)$ (red). It interacts with a parton j from h_2 with momentum fraction x_2 , factorised by $f_j(x_2)$, in the hard scattering process, described by the hard cross section $d\hat{\sigma}_{ij}$ (green). The final states of the hard scattering repeatedly decay, which can be approximated with a parton shower. Hadronisation (blue) occurs when the energy scale evolves below Λ_{QCD} .

independent parameters and is given by

$$d\Phi_N^{(d)}(\mathcal{I} \rightarrow \mathcal{F}) = (2\pi)^d \delta^{(d)}\left(\sum_{f \in \mathcal{F}} p_f - \sum_{i \in \mathcal{I}} p_i\right) \prod_{f \in \mathcal{F}} \frac{d^d p_f}{(2\pi)^{d-1}} \delta^{(+)}(p_f^2 - m_f^2), \quad (1.2.5)$$

where the first Dirac delta function fixes momentum conservation, and

$$\delta^{(+)}(p^2 - m^2) = \delta(p^2 - m^2) \theta(E), \quad (1.2.6)$$

is the on-shell condition for each external particle, with E as the temporal component of p . The Heaviside unit step function $\theta(E)$ selects the positive energy solution.

Thus, the quantity we need to compute is the ME. In processes consisting of D diagrams where D is a large number, it is more efficient to evaluate the D diagrams of the amplitude using the toolset built up in Sections 1.3 to 1.6 and then modulus square the resulting complex number to calculate the amplitude squared, rather than directly evaluating the D^2 terms of the squared amplitude.




Figure 1.5: A high-multiplicity collision event observed by the CMS detector at the LHC from collision data recorded in 2015 [66].

1.2.3 Beyond fixed order

Our discussion so far has assumed that the final states of the scattering are photons or partons. However, it is not hard partons that are measured at colliders. The initial and final state particles of the hard scattering process undergo additional radiation as they evolve to lower energy scales, creating a proliferation of extra particles.

The fixed-order prediction is constructed by integrating over the ME, so propagators (Section 1.3) give rise to logarithms. The presence of large scale separations in the kinematics thus leads to logarithmic corrections to the fixed-order prediction. Since soft and collinear states (Section 2.1) give rise to vanishing propagator momenta, these corrections can be significant in infrared (IR) regions of phase space for massless particles.

Treating these corrections to all orders in α_s is called resummation. It can be performed analytically [67, 68] in Soft-Collinear Effective Theory (SCET) [69–73], which systematically expands in powers of the logarithms, giving rise to the (next-to-)^kleading log ($N^k\text{LL}$) language, and kinematic variable, expressed as (next-to-)^kleading power ($N^k\text{LP}$). Resummation can also be done numerically below NLL through parton showers [74, 75], which is common with event generators [76, 77]. Matching logarithmic corrections to fixed-order results must be done carefully to avoid overcounting contributions [78].

All partons will eventually reach Λ_{QCD} and enter the soft regime, forming bound-state hadrons in a process known as hadronisation [79, 80]. For exclusive production (*i.e.* of a single hadron), for example, this can be described by fragmentation functions similarly to PDFs. These hadrons decay into child particles, which, along with further emissions such as photons, are what is actually measured. The parton shower and hadronisation processes within a hadronic simulation are depicted in Fig. 1.4.

The hard scattering products are then reconstructed from the measurements using jet clustering algorithms [81]; this is depicted in Fig. 1.5. Care must be taken that additional potentially unresolved emissions at higher orders (see Chapter 2) do not lead to ambiguities in the jets, which is called IR safety. Specifically, an IR safe observable is one for which the IR poles in the real- and virtual-type corrections occur in the same bins.

1.3 Amplitudes

We consider the scattering process of an initial collection of particles, denoted state $|i\rangle$, to a final state $|f\rangle$. The particles are characterised by properties such as their kinematics and particle type.

The probability of the initial state evolving to the final state is given by

$$P(i \rightarrow f) = |\langle f | S | i \rangle|^2, \quad (1.3.1)$$

where the scattering matrix S is a map between these states. We impose unitarity on S so that probability is conserved,

$$S^\dagger S = S S^\dagger = 1, \quad (1.3.2)$$

and explicitly separate the no-scattering case from the interaction, denoted by the transfer matrix T , as

$$S = 1 + iT. \quad (1.3.3)$$

We define the scattering amplitude \mathcal{A} [82] by writing the relation

$$(2\pi)^4 \delta^{(4)}(P_f - P_i) \mathcal{A}(i \rightarrow f) = \langle f | T | i \rangle, \quad (1.3.4)$$

where (P_i) P_f is the total momentum of the (initial) final state and the Dirac delta function imposes momentum conservation.

We calculate scattering amplitudes for n particles as a perturbative expansion in

a small dimensionless coupling parameter. For QCD, we use α_s ,

$$\mathcal{A}_n = \alpha_s^a \sum_{\ell=0}^L \alpha_s^\ell \mathcal{A}_n^{(\ell)} + \mathcal{O}(\alpha_s^{a+L+1}), \quad (1.3.5)$$

where a is the power in the coupling α_s of the leading order (LO) term and the $\mathcal{A}_n^{(\ell)}$ are sums over Feynman diagrams with n legs and ℓ loops. We truncate the expansion at some loop order L . The diagrams are composed of external legs, interaction vertices, and internal lines called propagators, with the possibility of closed propagators forming loops. We refer to diagrams with $\ell = 0$ as tree-level amplitudes or trees, $\ell = 1$ as one-loop amplitudes, and so on. Diagrams with $\ell \geq 1$ are called loop-level amplitudes. In the following, we are concerned with the efficient computation of the $\mathcal{A}_n^{(\ell)}$. Therefore, these order-by-order contributions are referred to as amplitudes, as well as the perturbative series which comprises the complete amplitude.

1.3.1 Properties

We can derive the mass dimension $[\mathcal{A}_n]$ of the $2 \rightarrow n - 2$ scattering amplitude by considering the cross section Eq. (1.2.2), which has dimensions of area, so mass dimension $[\hat{\sigma}] = -2$. The inverse flux factor Eq. (1.2.3) has mass dimension $[F] = -2$. The $(n - 2)$ -particle phase space, Eq. (1.2.5) with $N = n - 2$, has mass dimension $[\Phi_{n-2}^{(4)}] = 2n - 8$ in four dimensions¹. Thus, an amplitude of multiplicity (number of external legs) n must have mass dimension

$$[\mathcal{A}_n] = 4 - n. \quad (1.3.7)$$

Amplitudes are gauge-invariant objects. Thus, they obey the Ward identity. In QED, this means that the amplitude vanishes on replacing the polarisation vector of an external gauge boson with its momentum,

$$\mathcal{A}_n^{(\ell)}(p_1, \dots, p_n) =: \epsilon_\mu(p_i) \mathcal{A}_n^{(\ell)\mu}(p_1, \dots, p_n), \quad (p_i)_\mu \mathcal{A}_n^{(\ell)\mu}(p_1, \dots, p_n) = 0. \quad (1.3.8)$$

This also holds in QCD for amplitudes involving at most one longitudinal gluon.

Amplitudes exhibit two kinds of complexity [83]: algebraic, which arises from the increase in kinematic parameters at higher multiplicity; and analytic, which refers to the complicated functions that describe loops (Section 1.6), such as multiple

¹Recall the usual decomposition and change of variables for the Dirac delta function,

$$d^d p = d^{d-1} \vec{p} dE, \quad \delta(f(x)) = \sum_{x_0 \forall f(x_0)=0} \frac{\delta(x - x_0)}{f'(x_0)}. \quad (1.3.6)$$

polylogarithms and their elliptic generalisation [84–87].

1.3.2 Infinities

Tree-level diagrams have all internal line momenta fixed by the external momenta through momentum conservation at each vertex. However, in loop-level diagrams, there is an undetermined momentum flowing in each closed loop. Thus, each loop in an amplitude introduces an integral over an unconstrained momentum. These loop integrals can diverge—their value goes to infinity—in large-momentum regimes, where we call them ultraviolet (UV) divergences, and low-momentum regimes, which we call IR divergences. These points are also called singularities, and their treatment is called regularisation.

For instance, the dimensional regularisation (DR) method [88,89] prescribes that we regulate UV and IR divergences by performing an analytic continuation in the number of spacetime dimensions to

$$d = 4 - 2\epsilon. \quad (1.3.9)$$

This has the effect of modifying the four-dimensional phase-space integration measure to Eq. (1.2.5), and replacing each loop integration measure as

$$\alpha_s \frac{d^4 l_{[4]}}{(2\pi)^4} \rightarrow \alpha_s \mu^{2\epsilon} \frac{d^d l_{[d]}}{(2\pi)^d}, \quad (1.3.10)$$

where μ is the regularisation scale, an arbitrary scale introduced to fix the mass dimensions of the coupling. The loop integral is then expressed as a Laurent series in the dimensional regulator ϵ , with the singular behaviour cleanly extracted into poles in ϵ ; the $d = 4$ result is reproduced by taking $\epsilon \rightarrow 0$. DR has the advantage over simply introducing a UV cutoff that it preserves gauge invariance. Note that scaleless integrals¹ vanish in DR due to exact cancellation of the UV and IR singularities after analytic continuation.

We can also introduce the “spin dimension”, or quasidimension [84], d_s to regulate the contractions which arise from numerator algebra,

$$\eta^\mu{}_\mu = d_s, \quad (1.3.12)$$

¹Scaleless integrals are those that depend only on the regularisation scale μ . Since the integral then takes the form $I = \mu^a I'$ and μ is arbitrary, it must be that $I = 0$. For example,

$$\int \frac{d^d k}{(2\pi)^d} \frac{1}{k^2 (k+p)^2}, \quad p^2 = 0, \quad \int \frac{d^d k}{(2\pi)^d} \frac{1}{(k+q)^2}. \quad (1.3.11)$$

distinctly from the dimension of the loop integrals d . The conventional DR scheme is obtained by setting $d_s = d$. The 't Hooft-Veltman (tHV) scheme is given by setting $d_s = d$ only for parts of diagrams that can lead to divergences, and $d_s = 4$ otherwise. Thus, in tHV, external momenta are four-dimensional, while divergent loop-momenta are d -dimensional.

UV divergences arise due to our definition of the fields, couplings, and masses in the Lagrangian. These are “bare” quantities that must be related to physical quantities before our calculation can return a meaningful answer. The mechanism for this relationship is the scale dependence of parameters, as discussed for the running coupling in Section 1.1.2. The process is called renormalisation. We express the parameter rescalings as

$$Z_x = 1 + \delta_x, \quad (1.3.13)$$

and substitute them into the Lagrangian to introduce counterterms under some scheme, such as the modified minimal subtraction ($\overline{\text{MS}}$) scheme. Note that while it cancels their infinities, renormalisation is necessary regardless of the presence of UV divergent integrals.

Renormalised all-orders quantities are naturally independent of the scale μ_R at which renormalisation is carried out. However, truncated perturbative expansions in the coupling gain a residual scale dependence due to the omission of higher-order terms. It is common practice to set the regularisation scale μ equal to the renormalisation scale μ_R .

IR divergences are further discussed in Chapter 2.

1.4 Colour

In this section, the QCD gauge group is generalised to $\text{SU}(N_c)$, where N_c is the number of colours, to expose the group structure. For QCD, $N_c = 3$.

1.4.1 Colour decomposition

The number of Feynman diagrams contributing to an amplitude grows factorially with the number of legs. It is therefore essential to use techniques that reduce redundant intermediate calculation. One such method is colour decomposition, which exploits the property that QCD amplitudes can be factorised into colour and kinematic parts. Abstractly,

$$\mathcal{A}_n^{(\ell)} = \sum_i c_i^{(\ell)} A_{n,i}^{(\ell)} \quad (1.4.1)$$

where: $\mathcal{A}_n^{(\ell)}$ is the full amplitude at some loop order ℓ ; the colour factors $c_i^{(\ell)}$ contain the colour algebra; and the partial amplitudes $A_{n,i}^{(\ell)}$, also called colour-ordered or primitive amplitudes, contain the kinematics. The choice of decomposition is called a colour basis¹. In appropriate colour bases, the partial amplitudes are colour-ordered, meaning that they only receive contributions from diagrams with a particular cyclical ordering of the external partons, and remain gauge invariant. Thus, only adjacent legs may become collinear (Section 2.1).

Fundamental basis

The fundamental basis, also called the trace basis, is one such colour basis. It is constructed by taking strings of the $SU(N_c)$ generators in the fundamental representation t_{ij}^a , which are $N_c^2 - 1$ traceless hermitian $N_c \times N_c$ matrices [22, 90–92]. These matrices are normalised, in equation and diagram form, as

$$\text{tr}(t^a t^b) = T_F \delta^{ab}, \quad \text{loop diagram} = T_F \text{wavy line}. \quad (1.4.2)$$

The trace normalisation T_F , or Dynkin index, is $1/2$ for QCD, but will be left explicit in symbolic form. Note that the adjoint colour index is denoted by $a \in \{1, \dots, N_c^2 - 1\}$ and the fundamental colour index by $i, j \in \{1, \dots, N_c\}$.

Consider what can appear in a colour factor: a generic diagram will receive colour contributions from vertices,

$$\begin{array}{l} \begin{array}{c} i \\ \nearrow \\ \text{---} \\ \searrow \\ j \end{array} \text{---} a \quad \Big|_{\text{colour}} = t_{ij}^a, \\ \begin{array}{c} a \\ \text{---} \\ \text{---} \\ \text{---} \\ b \quad c \end{array} \quad \Big|_{\text{colour}} = f^{abc}, \quad \begin{array}{c} a \quad d \\ \text{---} \\ \text{---} \\ \text{---} \\ b \quad c \end{array} \quad \Big|_{\text{colour}} = f^{abe} f^{ecd}, \end{array} \quad (1.4.3)$$

and propagators,

$$a \text{---} b \Big|_{\text{colour}} = \delta^{ab}, \quad i \longleftarrow j \Big|_{\text{colour}} = \delta_{ij}. \quad (1.4.4)$$

¹Note that while the colour factors form a basis, the set of partial amplitudes may be linearly dependent.

We define the adjoint Casimir operator C_A ,

$$C_A = N_c, \quad f^{abc} f^{dcb} = C_A \delta^{ad}, \quad \text{[gluon loop]} = C_A \text{[gluon line]}, \quad (1.4.5)$$

and fundamental Casimir operator C_F ,

$$C_F = \frac{N_c^2 - 1}{2N_c}, \quad t_{ij}^a t_{jk}^a = C_F \delta_{ik}, \quad \text{[gluon self-energy]} = C_F \text{[gluon line]}. \quad (1.4.6)$$

Equation (1.1.11) allows us to recast the structure constants in terms of the generators as

$$f^{abc} = -\frac{i}{T_F} \text{tr} \left(t^a [t^b, t^c] \right),$$

$$\text{[3-gluon vertex]} = -\frac{i}{T_F} \left[\text{[triangle loop]} - \text{[triangle loop]} \right], \quad (1.4.7)$$

which allows all colour factors to be expressed in terms of the generators only. There also exists a symmetric tensor d^{abc} for the $SU(N_c)$ Lie algebra, defined by

$$\{t^a, t^b\} = \frac{1}{N_c} \delta^{ab} + d^{abc} t^c, \quad d^{abc} = \frac{1}{T_F} \text{tr} \left(t^a \{t^b, t^c\} \right), \quad (1.4.8)$$

although we will not make use of it here.

For example, consider n -gluon tree-level scattering amplitudes. After using the Fierz identity,

$$t_{ij}^a t_{kl}^a = T_F \left(\delta_{il} \delta_{kj} - \frac{1}{N_c} \delta_{ij} \delta_{kl} \right),$$

$$\text{[Fierz diagram]} = T_F \left[\text{[crossed lines]} - \frac{1}{N_c} \text{[self-energy]} \right], \quad (1.4.9)$$

which can be recognised as a completeness relation or proved via general tensor decomposition, and the cyclicity of the trace, the amplitude can be written in terms of $(n-1)!$ traces,

$$\mathcal{A}_n^{(0)} = \sum_{\sigma \in S_n / \mathbb{Z}_n} \text{tr} \left(\sigma(t^1, \dots, t^n) \right) A_n^{(0)}(\sigma(1, \dots, n)). \quad (1.4.10)$$

With S_n as the set of all permutations of n objects and \mathbb{Z}_n as the subset of cyclic permutations, the quotient S_n / \mathbb{Z}_n is the set of all non-cyclic permutations. This can be generated, for example, by fixing the first element and taking the cyclic permutations of the remaining elements.

The partial amplitudes also inherit a reflection identity from the antisymmetry of the colour-ordered Feynman rules, Eq. (1.4.3),

$$A_n^{(0)}(1, \dots, n) = (-1)^n A_n^{(0)}(n, \dots, 1), \quad (1.4.11)$$

which reduces the decomposition to $(n-1)!/2$ terms,

$$\mathcal{A}_n^{(0)} = \sum_{\sigma \in R_n} \lambda(\sigma(1, \dots, n)) A_n^{(0)}(\sigma(1, \dots, n)), \quad (1.4.12)$$

where

$$\lambda(1, \dots, n) = \text{tr}(t^1, \dots, t^n) + (-1)^n \text{tr}(t^n, \dots, t^1), \quad (1.4.13)$$

and the set R_n is the reflection-independent subset of S_n/\mathbb{Z}_n .

The partial amplitudes also obey the photon decoupling identity,

$$0 = A_n^{(0)}(1, 2, 3 \dots, n) + A_n^{(0)}(2, 1, 3 \dots, n) + \\ A_n^{(0)}(2, 3, 1 \dots, n) + \dots + A_n^{(0)}(2, 3 \dots, 1, n). \quad (1.4.14)$$

This arises because Eq. (1.4.10) is also valid for the gauge group $U(N_c) = \text{SU}(N_c) \times U(1)$, but the extra $U(1)$ gauge field, or “photon”, is colourless and therefore doesn’t couple to gluons, so any amplitude containing it must vanish. Thus, we insert a photon into Eq. (1.4.12) by setting one of the generators to be the generator of $U(1)$, which is proportional to the identity matrix, and gather terms to find Eq. (1.4.14). This means only $(n-2)!$ of the partial amplitudes are linearly independent (LI), so Eq. (1.4.12) is overcomplete.

The fundamental decomposition can be extended for general QCD amplitudes containing $q\bar{q}$ pairs and gluons, as well as loop-level partonic amplitudes.

Adjoint basis

Another colour basis is the adjoint basis [93, 94]. With the $\text{SU}(N_c)$ generator in the adjoint representation, Eq. (1.1.15c), the adjoint colour decomposition for the n -gluon tree-level amplitude is

$$\mathcal{A}_n^{(0)} = \sum_{\sigma \in S_{n-2}} [\sigma(F^2, \dots, F^{n-1})]_{1n} A_n^{(0)}(1, \sigma(2, \dots, n-1), n). \quad (1.4.15)$$

This decomposition can be derived from Eq. (1.4.10) using the Kleiss-Kuijff (KK) relation [95], which is consistent with the identities of Section 1.4.1 and describes

the linear relations between partial amplitudes. It can be stated as

$$A_n^{(0)}(1, \{\alpha\}, n, \{\beta\}) = (-1)^{n_\beta} \sum_{\sigma \in \text{OP}} A_n^{(0)}\left(1, \sigma\left(\{\alpha\}, \{\beta^T\}\right), n\right), \quad (1.4.16)$$

where

$$\{\alpha\} \cup \{\beta\} = \{2, \dots, n-1\}, \quad (1.4.17)$$

the length of the set $\{\beta\}$ is n_β , $\{\beta^T\}$ is $\{\beta\}$ with reversed ordering, and OP is the set of ordered permutations of $\{\alpha\} \cup \{\beta^T\}$, *i.e.* all orderings of the union that preserve the ordering of the elements within the sets $\{\alpha\}$ and $\{\beta^T\}$. The adjoint decomposition directly gives the $(n-2)!$ partial amplitudes.

The adjoint basis exists only for pure-gluon amplitudes.

Further bases

Many further approaches have been explored in the literature [96–100]. The Bern-Carrasco-Johansson relations project out $(n-3)!$ partial amplitudes [101, 102]. The multiplet basis provides a minimal basis [103–105]. To approximate two-loop processes, such as trijet production (Chapter 6), it is common to consider a colour expansion in $1/N_c^2$ (in the amplitude squared) and take only the leading-colour (LC) terms [106]. For calculations involving many particles, MC approximation techniques can be used which sample a reduced set of partial amplitudes in the colour sum [107–111]. The colour-flow basis [112, 113] treats gluons as $N_c \times N_c$ matrices $(A_\mu)^i_j$ rather than single-index fields A_μ^a such that the colour factor is a string of Kronecker delta functions in fundamental colour indices and the colour sum is over $(n-1)!$ terms; it requires the evaluation of fewer partial amplitudes in MC colour sums than other techniques. This has led to the development of the chirality-flow formalism [114–118] for the treatment of the kinematics, which builds on the spinor-helicity formalism (Section 1.5.4). The colour expansion approach has also been explored at high multiplicity [119].

1.4.2 Colour matrices

The calculation of cross sections requires the squared amplitude. Due to the phenomenon of colour confinement, colour is unobservable and so we average over initial state colours and sum over final state colours. Abstractly, we can introduce the colour matrix,

$$c_{ij}^{(\ell)} := \sum_{\text{colours}} c_i^{(\ell)\dagger} c_j^{(\ell)}, \quad (1.4.18)$$

which is symmetric and has integer-valued elements,

$$\mathcal{C}_{ij}^{(\ell)} = \mathcal{C}_{ji}^{(\ell)}, \quad \mathcal{C}_{ij}^{(\ell)} \in \mathbb{Z}, \quad (1.4.19)$$

to organise the computation as a colour sum, *c.f.* Eq. (1.4.1),

$$\sum_{\text{colours}} \left| \mathcal{A}_n^{(\ell)} \right|^2 = \sum_{i,j} A_{n,i}^{(\ell)\dagger} \mathcal{C}_{ij}^{(\ell)} A_{n,j}^{(\ell)}. \quad (1.4.20)$$

For example, consider the n -gluon tree-level amplitude. In the fundamental basis, using the decomposition Eq. (1.4.12), with

$$\begin{aligned} l_n &= 2 T_F^n N_c^{n-6} (N_c^2 - 1), \\ \gamma_1 &= N_c^4 - \gamma_2, \\ \gamma_2 &= 2 (N_c^2 - 3), \end{aligned} \quad (1.4.21)$$

the $n = 4$ colour matrix with corresponding partial amplitude vector is

$$\mathcal{C}_{ij}^{(0)} = l_4 \begin{pmatrix} \gamma_1 & \gamma_2 & \gamma_2 \\ \gamma_2 & \gamma_1 & \gamma_2 \\ \gamma_2 & \gamma_2 & \gamma_1 \end{pmatrix}_{ij}, \quad A_{4,i}^{(0)} = \begin{pmatrix} A_4^{(0)}(1, 2, 3, 4) \\ A_4^{(0)}(1, 2, 4, 3) \\ A_4^{(0)}(1, 3, 2, 4) \end{pmatrix}_i. \quad (1.4.22)$$

At $n = 5$, with

$$\begin{aligned} \gamma_3 &= N_c^4 - 4N_c^2 + 10, \\ \gamma_4 &= -2N_c^2 + 4, \end{aligned} \quad (1.4.23)$$

we have

$$\mathcal{C}_{ij}^{(0)} = l_5 \begin{pmatrix} \gamma_3 & \gamma_4 & \gamma_4 & 4 & 4 & -\gamma_4 & \gamma_4 & -4 & 4 & 0 & 4 & -\gamma_4 \\ \gamma_4 & \gamma_3 & 4 & -\gamma_4 & \gamma_4 & 4 & -4 & \gamma_4 & 0 & 4 & \gamma_4 & -4 \\ \gamma_4 & 4 & \gamma_3 & \gamma_4 & -\gamma_4 & 4 & 4 & 0 & -\gamma_4 & 4 & \gamma_4 & 4 \\ 4 & -\gamma_4 & \gamma_4 & \gamma_3 & 4 & \gamma_4 & \gamma_4 & -4 & -4 & -\gamma_4 & -4 & 0 \\ 4 & \gamma_4 & -\gamma_4 & 4 & \gamma_3 & \gamma_4 & 0 & 4 & 4 & -\gamma_4 & -4 & -\gamma_4 \\ -\gamma_4 & 4 & 4 & \gamma_4 & \gamma_4 & \gamma_3 & -4 & \gamma_4 & -\gamma_4 & -4 & 0 & 4 \\ \gamma_4 & -4 & 4 & \gamma_4 & 0 & -4 & \gamma_3 & \gamma_4 & \gamma_4 & 4 & -\gamma_4 & 4 \\ -4 & \gamma_4 & 0 & -4 & 4 & \gamma_4 & \gamma_4 & \gamma_3 & 4 & \gamma_4 & -4 & \gamma_4 \\ 4 & 0 & -\gamma_4 & -4 & 4 & -\gamma_4 & \gamma_4 & 4 & \gamma_3 & -\gamma_4 & 4 & \gamma_4 \\ 0 & 4 & 4 & -\gamma_4 & -\gamma_4 & -4 & 4 & \gamma_4 & -\gamma_4 & \gamma_3 & -\gamma_4 & -4 \\ 4 & \gamma_4 & \gamma_4 & -4 & -4 & 0 & -\gamma_4 & -4 & 4 & -\gamma_4 & \gamma_3 & \gamma_4 \\ -\gamma_4 & -4 & 4 & 0 & -\gamma_4 & 4 & 4 & \gamma_4 & \gamma_4 & -4 & \gamma_4 & \gamma_3 \end{pmatrix}_{ij}, \quad (1.4.24)$$

and

$$A_{5,i}^{(0)} = \begin{pmatrix} A_5^{(0)}(1, 2, 3, 4, 5) \\ A_5^{(0)}(1, 2, 3, 5, 4) \\ A_5^{(0)}(1, 2, 4, 3, 5) \\ A_5^{(0)}(1, 2, 4, 5, 3) \\ A_5^{(0)}(1, 2, 5, 3, 4) \\ A_5^{(0)}(1, 2, 5, 4, 3) \\ A_5^{(0)}(1, 3, 2, 4, 5) \\ A_5^{(0)}(1, 3, 2, 5, 4) \\ A_5^{(0)}(1, 3, 4, 2, 5) \\ A_5^{(0)}(1, 3, 5, 2, 4) \\ A_5^{(0)}(1, 4, 2, 3, 5) \\ A_5^{(0)}(1, 4, 3, 2, 5) \end{pmatrix}_i. \quad (1.4.25)$$

In the adjoint basis, with

$$k_n = \mathrm{T}_F^{n-2} N_c^{n-2} (N_c^2 - 1), \quad (1.4.26)$$

at $n = 4$,

$$\mathcal{C}_{ij}^{(0)} = k_4 \begin{pmatrix} 4 & 2 \\ 2 & 4 \end{pmatrix}_{ij}, \quad A_{4,i}^{(0)} = \begin{pmatrix} A_4^{(0)}(1, 2, 3, 4) \\ A_4^{(0)}(1, 3, 2, 4) \end{pmatrix}_i, \quad (1.4.27)$$

and at $n = 5$,

$$\mathcal{C}_{ij}^{(0)} = k_5 \begin{pmatrix} 8 & 4 & 4 & 2 & 2 & 0 \\ 4 & 8 & 2 & 0 & 4 & 2 \\ 4 & 2 & 8 & 4 & 0 & 2 \\ 2 & 0 & 4 & 8 & 2 & 4 \\ 2 & 4 & 0 & 2 & 8 & 4 \\ 0 & 2 & 2 & 4 & 4 & 8 \end{pmatrix}_{ij}, \quad A_{5,i}^{(0)} = \begin{pmatrix} A_5^{(0)}(1, 2, 3, 4, 5) \\ A_5^{(0)}(1, 2, 4, 3, 5) \\ A_5^{(0)}(1, 3, 2, 4, 5) \\ A_5^{(0)}(1, 3, 4, 2, 5) \\ A_5^{(0)}(1, 4, 2, 3, 5) \\ A_5^{(0)}(1, 4, 3, 2, 5) \end{pmatrix}_i. \quad (1.4.28)$$

The benefit of reducing the overcompleteness of the basis is evident in the reduced size of the colour matrix $\mathcal{C}_{ij}^{(0)}$ and partial amplitude vector $A_{n,i}^{(0)}$ for the adjoint basis compared to the fundamental. While it does not reduce the size of the colour sum—which becomes increasingly important at high multiplicity—relations between the partial amplitudes like the KK relation of Eq. (1.4.16) can, of course, instead be applied in the construction of the partial amplitude vector to optimise its computation in the fundamental basis.

1.5 Kinematics

Traditionally, the basic kinematic variables used for scattering amplitudes are four-momenta p_i^μ . Correspondingly in squared amplitudes, the momentum invariants and a pseudoscalar are used, introduced in Section 1.5.1. Using Weyl spinors (Section 1.5.3) to describe massless fermions leads to the spinor-helicity formalism (Section 1.5.4), which aims to express amplitudes more compactly. Momentum twistors (Section 1.5.7) provide another parametrisation that can lead to efficient expression.

Let us begin with some notation. The mostly-minus Minkowski metric tensor is used throughout,

$$\eta_{\mu\nu} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}_{\mu\nu}. \quad (1.5.1)$$

The $n \times n$ identity matrix is denoted by 1_n . The Pauli matrices are denoted by

$$\vec{\sigma} = (\sigma_1, \sigma_2, \sigma_3), \quad (1.5.2)$$

with

$$\sigma_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma_3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \quad (1.5.3)$$

We take all particles as massless, giving the on-shell constraint,

$$p^2 = 0. \quad (1.5.4)$$

Where quarks are considered, we are concerned only with the light quarks: u , d , c , s , and b . With the number of quark flavours denoted as N_f , this is the $N_f = 5$ regime. As we study hard scattering, this justifies the massless limit.

All particles are taken as outgoing, which reverses the helicity (the projection of spin onto the direction of momentum) of physically incoming particles. Thus, we consider $0 \rightarrow n$, or n -point, scattering; the physical $2 \rightarrow n - 2$ amplitude can be recovered by crossing symmetry. Conservation of four-momentum then conveniently reads as

$$\sum_{i=1}^n p_i^\mu = 0. \quad (1.5.5)$$

For example, we refer to five-gluon scattering as $5g$ in the all-outgoing case and $gg \rightarrow ggg$ in the physical case.

1.5.1 Degrees of freedom

We define the momentum invariants for n -point scattering as

$$s_{i\dots k} := (p_i + \dots + p_k)^2, \quad i, \dots, k \in \{1, \dots, n\}. \quad (1.5.6)$$

The two-index case s_{ij} is a generalisation to arbitrary multiplicity of the Mandelstam variables of $2 \rightarrow 2$ scattering. Note that they are symmetric,

$$s_{ij} = s_{ji}, \quad (1.5.7)$$

and for massless particles,

$$s_{ij} = 2 p_i \cdot p_j, \quad s_{ii} = 0. \quad (1.5.8)$$

Higher-index invariants decompose to double-index invariants and masses.

The parity-even kinematics can be described by a set of LI scalars. These can be constructed by contracting momenta, leading to a subset of the momentum invariants s_{ij} as a natural choice. Since they are symmetric, there are $n(n+1)/2$ distinct s_{ij} . For massless particles, we have n constraints—the diagonals s_{ii} vanish—taking us down to $n(n-1)/2$ ¹. There are a further n constraints,

$$\sum_{j=1}^n s_{ij} = 0 \quad \forall i \in \{1, \dots, n\}. \quad (1.5.9)$$

given by momentum conservation, Eq. (1.5.5). This gives the number of LI scalar products in d -dimensions as

$$\mathcal{D}^{(d)}(n) = \frac{n(n-3)}{2} \quad \forall n \leq (d+1). \quad (1.5.10)$$

We consider the momenta as d -dimensional because we are using DR. Giving some explicit cases,

$$\mathcal{D}^{(d)}(4) = 2, \quad \mathcal{D}^{(d)}(5) = 5, \quad \mathcal{D}^{(d)}(6) = 9. \quad (1.5.11)$$

To see why the formula Eq. (1.5.10) holds only for multiplicities $n \leq (d+1)$, consider the breakdown of the physical case $d = 4$ at $n = 6$. We have five “independent” momenta by momentum conservation, but as they are in a four-dimensional vector space, one must be dependent. Thus, we have

$$\mathcal{D}^{(4)}(6) = 8. \quad (1.5.12)$$

¹This can then be understood as n choose two combinations $\binom{n}{2}$ to generate all s_{ij} with $j > i$.

In the massless case, the scalars, s_{ij} of Eq. (1.5.8), form a Gram matrix of the external momenta when considered as an $m \times m$ matrix, where m gives the range of i and j as $i, j \in \{1, \dots, m\}$. If and only if the set of momenta chosen by the range of i and j are LI, then it is invertible and the determinant of the Gram matrix, or Gram determinant, is nonzero. Thus, the ($d = 4$, $n = 6$) Gram determinant with $m = 5$ vanishes, giving the additional constraint motivated before.

Considering also parity, which acts by flipping the sign of the spatial momentum components,

$$P : (p^0, \vec{p}) \longrightarrow (p^0, -\vec{p}) , \quad (1.5.13)$$

there must also be a parity-odd pseudoscalar invariant. In $d = 4$ with n particles, we can choose it as the parity-odd contraction of four LI four-momenta with the fully antisymmetric four-dimensional Levi-Civita symbol $\varepsilon_{\mu_1\mu_2\mu_3\mu_4}$. A common choice is

$$\text{tr}_5 := \text{tr}_5(p_1, p_2, p_3, p_4) , \quad (1.5.14)$$

where

$$\text{tr}_5(p_i, p_j, p_k, p_l) := 4i \varepsilon_{\mu_1\mu_2\mu_3\mu_4} p_i^{\mu_1} p_j^{\mu_2} p_k^{\mu_3} p_l^{\mu_4} . \quad (1.5.15)$$

We can also express this in terms of the spinor brackets defined in Section 1.5.4,

$$\text{tr}_5(p_i, p_j, p_k, p_l) = [i j] \langle j k \rangle [k l] \langle l i \rangle - \langle i j \rangle [j k] \langle k l \rangle [l i] . \quad (1.5.16)$$

Alternatively, this quantity can be expressed in terms of traces over gamma matrices,

$$\begin{aligned} \text{tr}_\pm(p_i, p_j, p_k, p_l) &= \frac{1}{2} \text{tr} \left[(1 \pm \gamma_5) \not{p}_i \not{p}_j \not{p}_k \not{p}_l \right] , \\ \text{tr}_+(p_i, p_j, p_k, p_l) &= [i j] \langle j k \rangle [k l] \langle l i \rangle , \\ \text{tr}_-(p_i, p_j, p_k, p_l) &= \langle i j \rangle [j k] \langle k l \rangle [l i] , \end{aligned} \quad (1.5.17)$$

by

$$\begin{aligned} \text{tr}_5(p_i, p_j, p_k, p_l) &= \text{tr}_+(p_i, p_j, p_k, p_l) - \text{tr}_-(p_i, p_j, p_k, p_l) \\ &= \text{tr} \left(\gamma^5 \not{p}_i \not{p}_j \not{p}_k \not{p}_l \right) . \end{aligned} \quad (1.5.18)$$

Note that for $n = 3$, since we do not have sufficient LI four-momenta, the pseudoscalar product Eq. (1.5.15) vanishes and the kinematics are purely parity-even. Since the product of two pseudoscalars is a scalar, all contractions of a pseudoscalar with a scalar can be expressed as one pseudoscalar variable, such as tr_5 , multiplied by some algebraic function of scalar variables, such as s_{ij} .

The square of tr_5 can be expressed in terms of the scalars through the Gram

determinant with

$$\mathrm{tr}_5^2 = \Delta := \det(s_{ij})_{i,j \in \{1,2,3,4\}}, \quad (1.5.19)$$

which is a degree-4 polynomial in the s_{ij} . The pseudoscalar invariant tr_5 therefore introduces an algebraic dependence on the kinematics, since

$$\mathrm{tr}_5 = \pm\sqrt{\Delta}. \quad (1.5.20)$$

We emphasise that the sign of tr_5 changes under parity conjugation, Eq. (1.5.13), and under odd-signature permutations of the external momenta.

For further reading, see Ref. [120] and Section 2 of Ref. [121].

1.5.2 Representations

The Lorentz group¹ is the Lie group $\mathrm{SO}(1,3)$. The algebra of this group, $\mathfrak{so}(1,3)$, maps to two commuting copies of the algebra of $\mathrm{SU}(2)$, $\mathfrak{su}(2)$, so its representations can be labelled by the chiral doublet (a,b) where a and b are the eigenvalues of the Casimir operators of the two $\mathfrak{su}(2)$ algebras.

Four-momenta p^μ , with Lorentz indices $\mu \in \{0,1,2,3\}$, are Lorentz vectors that transform under the $(1/2,1/2)$ representation of the Lorentz group. In the spinor-helicity formalism, spinors² are instead used as the basic kinematic variable. In general, Dirac spinors ψ^α , with Dirac spinor indices $\alpha \in \{1,2,3,4\}$, might be used, which solve the Dirac equation and lie in the $(1/2,0) \oplus (0,1/2)$ representation. For massless vectors, however, the Dirac spinor decomposes into two Weyl spinors and we use these instead. The first is the left-handed Weyl spinor $\lambda^{\dot{\alpha}}$, also called the holomorphic spinor, lying in the $(1/2,0)$ representation with left-handed Weyl spinor indices $\dot{\alpha} \in \{\dot{1}, \dot{2}\}$. The other is the right-handed Weyl spinor $\tilde{\lambda}^\alpha$, or antiholomorphic spinor, lying in the $(0,1/2)$ representation with right-handed Weyl spinor indices $\alpha \in \{1,2\}$. There also exist negative energy solutions to the Weyl equations, but for massless particles they are equal to the positive solutions and hence can be neglected.

1.5.3 Weyl spinors

In analogy to the metric tensor in a Lorentz vector space, the raising and lowering of indices in the Weyl spinor spaces is handled by the two-dimensional Levi-Civita

¹Strictly, ‘‘Lorentz group’’ is used here to refer to the identity connected component of the $\mathrm{O}(1,3)$ group.

²Strictly, spinors lie in representations of the spin group $\mathrm{SL}(2,\mathbb{C})$ which is a homomorphism of the Lorentz group.

tensor ε (choosing $\varepsilon^{12} = \varepsilon^{\dot{1}\dot{2}} = 1$ such that $\varepsilon = i\sigma_2$), for example,

$$\lambda_{\dot{\alpha}} = \varepsilon_{\dot{\alpha}\dot{\beta}} \lambda^{\dot{\beta}}, \quad \tilde{\lambda}_{\alpha} = \varepsilon_{\alpha\beta} \tilde{\lambda}^{\beta}. \quad (1.5.21)$$

Translations between Lorentz and spinor indices use the sigma matrices,

$$(\sigma^{\mu})_{\alpha\dot{\alpha}} := (1_2, \vec{\sigma}), \quad (\bar{\sigma}^{\mu})^{\dot{\alpha}\alpha} := (1_2, -\vec{\sigma}), \quad (1.5.22)$$

noting that

$$(\bar{\sigma}_{\mu})^{\dot{\alpha}\alpha} = \varepsilon^{\alpha\beta} \varepsilon^{\dot{\alpha}\dot{\beta}} (\sigma_{\mu})_{\dot{\beta}\beta}. \quad (1.5.23)$$

The sigma matrices are analogous to the gamma matrices of Dirac spinor algebra.

The four-momentum Lorentz vector p^{μ} is related to a rank-two spinor (*i.e.* with two indices), or bispinor, by

$$p^{\dot{\alpha}\alpha} = (\bar{\sigma}_{\mu})^{\dot{\alpha}\alpha} p^{\mu}. \quad (1.5.24)$$

This bispinor is the Weyl spinor analogy to a slashed momentum in the Feynman slash notation, Eq. (1.1.3). Expressions in Dirac spinors for massless fermions with gamma matrices can always be decomposed into Weyl spinors and sigma matrices by choosing the Weyl, or chiral, representation of the gamma matrices,

$$\gamma^{\mu} = \begin{pmatrix} 0 & \sigma^{\mu} \\ \bar{\sigma}^{\mu} & 0 \end{pmatrix}, \quad \gamma^5 = \begin{pmatrix} -1_2 & 0 \\ 0 & 1_2 \end{pmatrix}. \quad (1.5.25)$$

Similarly, the massless Dirac equation,

$$\not{p} \psi(p) = 0, \quad (1.5.26)$$

decomposes into the left- and right-handed Weyl equations,

$$p_{\alpha\dot{\alpha}} \lambda^{\dot{\alpha}} = 0, \quad p^{\dot{\alpha}\alpha} \tilde{\lambda}_{\alpha} = 0. \quad (1.5.27)$$

The definition of the sigma matrices leads to the Fierz identities,

$$(\sigma^{\mu})_{\alpha\dot{\alpha}} (\bar{\sigma}_{\mu})^{\dot{\beta}\beta} = 2 \delta_{\alpha}^{\beta} \delta_{\dot{\alpha}}^{\dot{\beta}}, \quad \text{tr}(\sigma^{\mu} \bar{\sigma}^{\nu}) = 2 \eta^{\mu\nu}, \quad (1.5.28)$$

which can be used to express the on-shell constraint Eq. (1.5.4) as

$$\frac{1}{2} p^{\dot{\alpha}\alpha} p^{\dot{\beta}\beta} \varepsilon_{\dot{\alpha}\dot{\beta}} \varepsilon_{\alpha\beta} = \det(p^{\dot{\alpha}\alpha}) = 0. \quad (1.5.29)$$

Since the determinant vanishes, the 2×2 matrix $p^{\dot{\alpha}\alpha}$ is rank one (*i.e.* has a single

LI row or column). It may be written in bispinor form as

$$p^{\dot{\alpha}\alpha} = \lambda^{\dot{\alpha}} \tilde{\lambda}^{\alpha}. \quad (1.5.30)$$

For real momentum, this means that

$$\lambda^{\dot{\alpha}} = \left(\tilde{\lambda}^{\alpha} \right)^*, \quad (1.5.31)$$

or in other words, complex conjugation is equivalent to a chirality flip. If the momentum is complex, then the spinors are unrelated. The bispinor form also demonstrates that the momentum is invariant under little group transformations¹ [122],

$$\left(\lambda, \tilde{\lambda} \right) \rightarrow \left(\phi \lambda, \frac{\tilde{\lambda}}{\phi} \right), \quad (1.5.32)$$

for a complex phase $\phi \in \mathbb{C}$, so a four-momentum does not map uniquely to a spinor pair.

1.5.4 Spinor-helicity formalism

In the spinor-helicity formalism, the scattering amplitude of n massless particles is a function of the set of Weyl spinors $\{\lambda_i, \tilde{\lambda}_i\}$ with particle label $i \in \{1, \dots, n\}$. Since they describe massless fermions, Weyl spinors with left-handed chirality λ have negative helicity, and similarly the right-handed spinors $\tilde{\lambda}$ have positive helicity.

A bracket notation is used to denote the Weyl spinors, defined by the spinor products,

$$\begin{aligned} [i j] &:= (\lambda_i)_{\dot{\alpha}} \lambda_j^{\dot{\alpha}} = (\lambda_i)_{\dot{\alpha}} \varepsilon^{\dot{\alpha}\dot{\beta}} (\lambda_j)_{\dot{\beta}}, \\ \langle i j \rangle &:= \tilde{\lambda}_i^{\alpha} (\tilde{\lambda}_j)_{\alpha} = (\tilde{\lambda}_i)_{\alpha} \varepsilon^{\alpha\beta} (\tilde{\lambda}_j)_{\beta}, \end{aligned} \quad (1.5.33)$$

such that,

$$\begin{aligned} |i] &= \lambda_i^{\dot{\alpha}}, & [i] &= (\lambda_i)_{\dot{\alpha}}, \\ |i\rangle &= (\tilde{\lambda}_i)_{\alpha}, & \langle i| &= \tilde{\lambda}_i^{\alpha}. \end{aligned} \quad (1.5.34)$$

Contractions with matrices are denoted by

$$[i \sigma^{\mu} j\rangle := (\lambda_i)_{\dot{\alpha}} (\sigma^{\mu})^{\dot{\alpha}\alpha} (\tilde{\lambda}_j)_{\alpha}, \quad (1.5.35)$$

¹The little group is the group of transformations for which the four-momentum of an on-shell particle is invariant.

and contractions with bispinors as

$$\begin{aligned} [i j k] &:= p_j^\mu [i \sigma_\mu k] = (\lambda_i)_{\dot{\alpha}} (p_j)^{\dot{\alpha}\alpha} (\tilde{\lambda}_k)_\alpha, \\ \langle i j k \rangle &:= p_j^\mu \langle i \bar{\sigma}_\mu k \rangle = \tilde{\lambda}_i^\alpha (p_j)_{\alpha\dot{\alpha}} \lambda_k^{\dot{\alpha}}. \end{aligned} \quad (1.5.36)$$

The manipulation of expressions involving these square and angle bracket spinors mainly involves the use of a number of identities. Some common ones are listed here:

- Momentum invariants

$$s_{ij} = \langle i j \rangle [j i]. \quad (1.5.37)$$

- Antisymmetry

$$\langle j i \rangle = -\langle i j \rangle, \quad [j i] = -[i j], \quad \langle i i \rangle = [i i] = 0. \quad (1.5.38)$$

- Projection operators

$$|i\rangle[i] = (p_i)_{\alpha\dot{\alpha}}, \quad |i\rangle\langle i| = (p_i)^{\dot{\alpha}\alpha}. \quad (1.5.39)$$

- Gordon identity

$$[i \sigma^\mu i] = 2 p_i^\mu. \quad (1.5.40)$$

- Fierz identity

$$[i \sigma^\mu j][k \sigma_\mu l] = 2 [i k] \langle l j \rangle. \quad (1.5.41)$$

- Charge conjugation

$$[i \sigma^\mu j] = \langle j \bar{\sigma}^\mu i \rangle. \quad (1.5.42)$$

- Complex conjugation (for real momenta $p_i^\mu \in \mathbb{R}$ only)

$$\langle i j \rangle^* = \begin{cases} [i j] & \text{if } \text{sign}(p_i^0) = \text{sign}(p_j^0), \\ -[i j] & \text{otherwise.} \end{cases} \quad (1.5.43)$$

- Schouten identities¹

$$\langle i j \rangle \langle k l \rangle = \langle i k \rangle \langle j l \rangle + \langle i l \rangle \langle k j \rangle, \quad [i j][k l] = [i k][j l] + [i l][k j]. \quad (1.5.44)$$

¹This follows from the fact that a two-component object can be written as the linear combination of two others.

- Momentum conservation for an n -point amplitude

$$\sum_{j=1}^n [i j] \langle j k \rangle = 0 \quad \forall i, k \in \{1, \dots, n\} . \quad (1.5.45)$$

To express all the kinematics of a generic amplitude in terms of the spinor variables, the polarisation vectors of bosons also require a spinor representation. This is achieved for massless gauge bosons of definite helicity ± 1 by

$$\varepsilon_{\mu}^{-}(p_i, p_r) = -\frac{[r \sigma_{\mu} i]}{\sqrt{2} [r i]}, \quad \varepsilon_{\mu}^{+}(p_i, p_r) = \frac{\langle r \bar{\sigma}_{\mu} i \rangle}{\sqrt{2} \langle r i \rangle}, \quad (1.5.46)$$

where the superscript \pm refers to the helicity of the boson, p_i is the momentum of the boson, and p_r is a reference momentum due to gauge invariance. We define

$$\varepsilon_i^{\lambda}(q) := \varepsilon^{\lambda}(p_i, q) \quad (1.5.47)$$

for convenience. The amplitude is independent of the choice of reference momentum; judicious choice of its value can greatly simplify algebra. The Weyl equation, *c.f.* Eq. (1.5.27),

$$[r i i] = 0, \quad (1.5.48)$$

implies that the polarisation vector is transverse to its momentum for any reference momentum,

$$\varepsilon_i^{\pm}(q) \cdot p_i = 0. \quad (1.5.49)$$

Again, there are some useful identities to deal with these expressions:

- Complex conjugation

$$(\varepsilon_i^{\mp}(q))^* = \varepsilon_i^{\pm}(q) . \quad (1.5.50)$$

- Normalisation

$$\varepsilon_i^{\pm}(q) \cdot (\varepsilon_i^{\pm}(q))^* = -1. \quad (1.5.51)$$

- Orthogonality

$$\varepsilon_i^{\pm}(q) \cdot (\varepsilon_i^{\mp}(q))^* = 0. \quad (1.5.52)$$

- Completeness relation

$$\sum_{\lambda \in \{-, +\}} \varepsilon_\mu^\lambda(p, q) \left(\varepsilon_\nu^\lambda(p, q) \right)^* = -\eta_{\mu\nu} + \frac{p_\mu q_\nu + p_\nu q_\mu}{q \cdot p}. \quad (1.5.53)$$

- Cancellation identities

$$\varepsilon_i^\pm(q) \cdot q = 0, \quad (1.5.54a)$$

$$\varepsilon_i^\pm(q) \cdot \varepsilon_j^\pm(q) = 0, \quad (1.5.54b)$$

$$\varepsilon_i^\pm(p_j) \cdot \varepsilon_j^\mp(q) = 0, \quad (1.5.54c)$$

$$\not\epsilon_i^+(p_j) |j\rangle = \not\epsilon_i^-(p_j) |j\rangle = 0, \quad (1.5.54d)$$

$$[j | \not\epsilon_i^-(p_j) = \langle j | \not\epsilon_i^+(p_j) = 0. \quad (1.5.54e)$$

With these identities, it is possible to express and simplify any amplitude with massless fermions and vector bosons as the external states in the spinor-helicity formalism.

Recall the little group scaling of the Weyl fermions, Eq. (1.5.32), or in the new notation,

$$\begin{aligned} |i\rangle &\rightarrow \phi |i\rangle, & |i\rangle &\rightarrow \phi^{-1} |i\rangle, \\ [i] &\rightarrow \phi [i], & \langle i| &\rightarrow \phi^{-1} \langle i|. \end{aligned} \quad (1.5.55)$$

We now scale Eq. (1.5.46) to find for the bosons that

$$\varepsilon_i^\pm \rightarrow \phi^{\mp 2} \varepsilon_i^\pm. \quad (1.5.56)$$

Thus, all particles transform as ϕ^{-2h} , where h is the helicity of the particle,

$$h = \begin{cases} -\frac{1}{2} & \text{for } |i\rangle, \\ \frac{1}{2} & \text{for } |i\rangle, \\ \pm 1 & \text{for } \varepsilon^\pm. \end{cases} \quad (1.5.57)$$

Therefore, an amplitude of n massless particles scales according to the helicity of each particle,

$$\prod_{i=1}^n \phi_i^{-2h_i}. \quad (1.5.58)$$

These scaling factors are also referred to as phase or helicity weights.

1.5.5 Helicity amplitudes

Treating spin on the assumption that the collider uses an unpolarised input and does not measure the polarisation of the product states, we construct the full amplitude as a sum over helicity amplitudes,

$$\mathcal{A}_n^{(\ell)} = \sum_{h_1, \dots, h_n} \mathcal{A}_n^{(\ell)}(1^{h_1}, \dots, n^{h_n}), \quad (1.5.59)$$

where $h_i \in \{+, -\}$. For tree-level helicity amplitudes, the polarisation vector cancellation identity Eq. (1.5.54b) causes all amplitudes of helicity signature¹ $(+, -) = (0, n)$ and $(1, n-1)$, and their parity conjugates, to vanish,

$$A_n^{(0)}(1^\pm, \dots, n^\pm) = A_n^{(0)}(1^\pm, \dots, i^\mp, \dots, n^\pm) = 0. \quad (1.5.60)$$

Therefore, amplitudes with signature $(2, n-2)$ receive the name maximally-helicity-violating (MHV) amplitudes, and their conjugates are called anti-maximally-helicity-violating ($\overline{\text{MHV}}$) amplitudes. Signature $(3, n-3)$ amplitudes are called next-to-maximally-helicity-violating (NMHV) amplitudes, and so on.

To calculate pure-gluon tree-level MHV amplitudes, the tools we have now built up come into good use. The number of diagrams to evaluate is reduced using colour decomposition and relations between partial amplitudes. Others vanish due to judicious choice of the reference momenta. Finally, use of the identities of Section 1.5.4 allow the result to be collected in a single term. A remarkably simple expression is found as the result: the Parke-Taylor amplitudes [123, 124], which have the form

$$\begin{aligned} A_n^{(0), \text{MHV}}(1^+, \dots, j^-, \dots, k^-, \dots, n^+) &= i \frac{\langle j k \rangle^4}{\prod_{i=1}^n \langle i | i+1 \rangle}, \\ A_n^{(0), \overline{\text{MHV}}}(1^-, \dots, j^+, \dots, k^+, \dots, n^-) &= i \frac{[j k]^4}{\prod_{i=1}^n [i | i+1]}, \end{aligned} \quad (1.5.61)$$

with the indices i defined modulo n .

Massless three-particle scattering presents a special case. By momentum conservation Eq. (1.5.5), $s_{12} = 0$, and so by Eq. (1.5.37), either $\langle 12 \rangle$ or $[12]$ must vanish. Furthermore, by Eqs. (1.5.27) and (1.5.45), we find either all square brackets or all angle brackets vanish,

$$\langle 12 \rangle = \langle 23 \rangle = \langle 13 \rangle = 0 \quad \text{or} \quad [12] = [23] = [13] = 0. \quad (1.5.62)$$

Since a nonzero three-particle amplitude must therefore depend only on angle or square brackets, and for real momenta, these are each other's complex conjugates by

¹Helicity signature (a, b) refers to an $(a+b)$ -particle amplitude that has a positive helicity legs and b negative helicity legs.

Eq. (1.5.43), such amplitudes do not exist for physical kinematics. However, they are nonvanishing with complex momenta, with values given by Eq. (1.5.61).

We use curly brackets to denote the exponents of the little group scalings Eq. (1.5.58) for each leg of helicity amplitudes, or in general any object carrying helicity weights. For example, consider the four-gluon helicity amplitude,

$$\{\mathcal{A}_4(1_g^+, 2_g^+, 3_g^-, 4_g^-)\} = \{-2, -2, 2, 2\}. \quad (1.5.63)$$

Beyond tree level in the amplitude—or equivalently LO in the amplitude squared, assuming the process is not loop induced, as discussed in Section 1.8.1—the helicity configurations of Eq. (1.5.60) are finite. These are referred to as ultra-helicity-violating (UHV) amplitudes. At next-to-leading order (NLO), the virtual correction is the interference of one-loop amplitudes and tree-level amplitudes. Therefore, UHV amplitudes first appear at next-to-next-to-leading order (NNLO) in the double-virtual (VV) corrections as one-loop squared amplitudes.

1.5.6 Recursion

In Section 1.4.1, colour decomposition was used to reduce the number of Feynman diagrams required to calculate an amplitude. In this spirit, it is also possible to construct higher-multiplicity amplitudes recursively from lower-multiplicity amplitudes. This is a useful technique in the numerical evaluation of high-multiplicity amplitudes. One such method is Berends-Giele off-shell recursion for partial amplitudes [124]. A partial amplitude is expressed as

$$A_n^{(0)}(1, \dots, n) = \varepsilon_\mu(p_n) J^\mu(1, \dots, n-1), \quad (1.5.64)$$

where J^μ is called the off-shell current. In the case of gluon amplitudes, it is defined recursively as

$$J^\mu(1, \dots, n) = \frac{-i}{s_{1,n}} \left[\sum_{i=1}^{n-1} V_3^{\mu\nu\rho}(p_{1,i}, p_{i+1,n}) J_\nu(1, \dots, i) J_\rho(i+1, \dots, n) + \sum_{i=1}^{n-2} \sum_{j=i+1}^{n-1} V_4^{\mu\nu\rho\sigma} J_\nu(1, \dots, i) J_\rho(i+1, \dots, j) J_\sigma(j+1, \dots, n) \right], \quad (1.5.65)$$

with

$$p_{i,k} := p_i + \dots + p_k, \quad s_{i,k} := p_{i,k}^2. \quad (1.5.66)$$

The gluon vertices $V_3^{\mu\nu\rho}$ and $V_4^{\mu\nu\rho\sigma}$ are given in $R_{\xi=1}$ gauge by the colour-ordered Feynman rules,

$$\begin{aligned} V_3^{\mu\nu\rho}(p, q) &= i(\eta^{\nu\rho}(p - q)^\mu + 2\eta^{\rho\mu}q^\nu - 2\eta^{\mu\nu}p^\rho), \\ V_4^{\mu\nu\rho\sigma} &= i(2\eta^{\mu\rho}\eta^{\nu\sigma} - \eta^{\mu\nu}\eta^{\rho\sigma} - \eta^{\mu\sigma}\eta^{\nu\rho}). \end{aligned} \quad (1.5.67)$$

Pictorially, Eq. (1.5.65) can be stated as

$$\begin{array}{c} \begin{array}{c} 1 \\ \vdots \\ \mu \text{ --- } \textcircled{J} \\ \vdots \\ n \end{array} \end{array} = \frac{-i}{s_{1,n}} \left(\sum_i \begin{array}{c} 1 \\ \vdots \\ \textcircled{J} \text{ --- } i \\ \vdots \\ \textcircled{J} \text{ --- } i+1 \\ \vdots \\ n \end{array} + \sum_{i,j} \begin{array}{c} 1 \\ \vdots \\ \textcircled{J} \text{ --- } i \\ \vdots \\ \textcircled{J} \text{ --- } j \\ \vdots \\ n \end{array} \right). \quad (1.5.68)$$

The off-shell current J_μ has one off-shell leg, labelled by μ , and is therefore gauge dependent. It obeys the photon decoupling equation Eq. (1.4.14), the reflection identity Eq. (1.4.11) with an extra sign, and current conservation,

$$p_{1,n}^\mu J_\mu(1, \dots, n) = 0. \quad (1.5.69)$$

An amplitude is recovered on contraction of the off-shell current with a polarisation vector, as in Eq. (1.5.64). Off-shell recursion can be extended to additionally include quarks [125].

An alternative method is Britto-Cachazo-Feng-Witten (BCFW) on-shell recursion [126, 127], which can be advantageous for helicity expressions, although it requires the introduction of complex momenta. Recursive techniques allow common building blocks of an amplitude to be reused so that the number of independent diagram evaluations is reduced to polynomial growth in the number of legs [128].

1.5.7 Momentum twistor variables

Our methods so far have cast amplitudes as analytical expressions of the scalars s_{ij} , the pseudoscalar tr_5 , and the spinor products $\langle ij \rangle$ and $[ij]$. One advantage of these variables is that the poles in the kinematics (see Chapter 2) are manifest in expressions [129]. While these variables are highly linearly dependent, it is possible to reduce them to an LI set. However, it can be the case that using a particular choice of overcomplete set leads to more compact expressions. Automating the

reduction to a set that gives the shortest expressions is very challenging. In addition, recall that tr_5 introduces algebraic dependence on the kinematics with a square root, Eq. (1.5.20).

An alternative representation is provided by momentum twistor parametrisation [130, 131], in which n -point kinematics are represented by a momentum twistor matrix,

$$Z = \left(Z_1 \cdots Z_n \right). \quad (1.5.70)$$

Each momentum twistor,

$$Z_i = \begin{pmatrix} \lambda_i \\ \mu_i \end{pmatrix}, \quad (1.5.71)$$

has the first two components as those of the negative-helicity Weyl spinor λ_i , while the last two components μ_i are defined through their relation to the positive-helicity Weyl spinor $\tilde{\lambda}_i$,

$$\tilde{\lambda}_i = \frac{\langle i|i+1\rangle \mu_{i-1} + \langle i+1|i-1\rangle \mu_i + \langle i-1|i\rangle \mu_{i+1}}{\langle i|i+1\rangle \langle i-1|i\rangle}, \quad (1.5.72)$$

with the indices defined modulo n . Thus, Z is an $4 \times n$ matrix and so has $4n$ momentum twistor components.

The momentum twistor matrix Z enjoys Poincaré symmetry, which has ten generators. The corresponding physical momentum p_i is invariant under a $U(1)$ scaling of the components of a momentum twistor,

$$\lambda_i \rightarrow \phi_i \lambda_i, \quad \mu_i \rightarrow \phi_i \mu_i, \quad \phi_i \in U(1), \quad (1.5.73)$$

which arises from the little group invariance Eq. (1.5.32). Consequently, each momentum twistor is defined projectively,

$$p_i \rightarrow p_i \quad \text{under} \quad Z_i \rightarrow \phi_i Z_i. \quad (1.5.74)$$

Therefore, Z has the number of independent components,

$$\mathcal{D}^{(4)}(n) = 4n - 10 - n = 3n - 10. \quad (1.5.75)$$

This agrees with our ($d = 4$, $n \leq 6$) parity-even analysis in Section 1.5.1. The choice of these free components, which we refer to as momentum twistor variables (MTVs), is not unique. They are also referred to as momentum twistor coordinates.

MTVs are unconstrained by construction, in that both massless on-shell constraints Eq. (1.5.4) and momentum conservation Eq. (1.5.5) are automatically satisfied. Hence, the Schouten identities Eq. (1.5.44) are trivially satisfied when written

in terms of MTVs. Identifying relations between functions of MTVs is therefore easy, and the automated simplification of expressions is unambiguous. It is also straightforward to convert them back to traditional kinematic variables.

Another advantage of MTVs is that they can provide a rational parametrisation of the kinematics: the spinors λ and $\tilde{\lambda}$, and momenta p , and therefore the spinor products, scalars s_{ij} , and pseudoscalar tr_5 , are all rational functions of the MTVs. This is necessary for the use of finite field (FF) techniques (Section 1.7), as we will see in Section 4.3.

A disadvantage of MTVs is that they do not contain the helicity phase information, which is lost when translating from spinor products. However, the correct phase weights can easily be restored to an MTV expression by multiplying by a suitable phase factor.

MTVs are invaluable tools to provide rational kinematics and obtain compact expressions in traditional variables, as well as providing efficient expressions themselves. The geometric origin of momentum twistors is discussed in Sections 5.3 and 5.4 of Ref. [20]. We discuss an explicit parametrisation in Section 4.4.

1.6 On-shell methods and integral reduction

The implementation of on-shell methods and integral reduction led to an “NLO revolution” in the automated computation of one-loop amplitudes [63, 132–139]. This technology has recently been applied to automate two-loop amplitudes [140, 141], but evaluation is relatively costly. This indicates the need for a mixed analytical and numerical approach at NNLO, such as reconstruction of analytic expressions (Section 1.7). For further reading on the techniques discussed in this section, see Refs. [142–146].

1.6.1 Loops

An L -loop partial amplitude can be written as a sum over all contributing L -loop diagrams,

$$A_n^{(L)} = \sum_j \int \left(\prod_{\ell=1}^L \frac{d^d k_\ell}{(2\pi)^d} \right) \partial A_{n,j}^{(L)}, \quad (1.6.1)$$

where k_ℓ are the loop-momenta, and the diagram integrands take the general form,

$$\partial A_{n,j}^{(L)} = S_j \frac{N_j}{\prod_{\alpha_j} D_{\alpha_j}}, \quad (1.6.2)$$

where S_j is a symmetry factor, α_j index the propagators, D_{α_j} are inverse propagators, and N_j are polynomials in Lorentz-invariant contractions of loop-momenta, external momenta, and polarisation vectors.

After regulating the divergences (Section 1.3.2), loop integrals are typically not expressible as rational functions. A common analytic structure in massless amplitudes are polylogarithms $\text{Li}_i(z)$ which are defined iteratively by

$$\text{Li}_i(z) = \int_0^z dy \frac{\text{Li}_{i-1}(y)}{y}, \quad \text{Li}_1(z) = -\ln(1-z), \quad (1.6.3)$$

where $\ln(z)$ is the natural logarithm (with base e). Such functions are referred to as transcendental functions, and have a transcendental weight, or degree of transcendentality, given by the dimension of the integral(s). Hence, $\ln(z)$ and π have transcendental weight one, since

$$\ln(z) = \int_1^z \frac{dt}{t}, \quad i\pi = \ln(-1), \quad (1.6.4)$$

and $\text{Li}_i(z)$ has transcendental weight i .

1.6.2 Unitarity cuts

Let us take some massless one-loop amplitude and choose the loop-momentum k in each diagram such that we can rewrite Eq. (1.6.1) under one integral,

$$A_n^{(1)} = \int \frac{d^d k}{(2\pi)^d} \sum_j \partial A_{n,j}^{(1)}, \quad D_{\alpha_j} = \left(k - \sum_{\beta_j} p_{\beta_j} \right)^2, \quad \alpha_j \in \{1, \dots, n\}, \quad (1.6.5)$$

where p_i are the external momenta. We can apply cut conditions on the loop-momentum where we impose that certain loop propagators α_j have zero momentum. In other words, we take these lines on-shell, which we call taking unitarity cuts. Then, for this particular subplane of loop-momentum, the loop amplitude factorises as a product of on-shell tree-level amplitudes.

We will now arrive at this relationship from unitarity of the S -matrix. We insert Eq. (1.3.3) into Eq. (1.3.2),

$$\begin{aligned} T^\dagger T &= i (T^\dagger - T) \\ &= 2 \text{Im}(T), \end{aligned} \quad (1.6.6)$$

and sandwich with states,

$$2 \text{Im}(\langle f | T | i \rangle) = \sum_a \langle f | T^\dagger | a \rangle \langle a | T | i \rangle, \quad (1.6.7)$$

where the insertion of identity involves a sum over the degrees of freedom of the state a . We can write this in terms of amplitudes using Eq. (1.3.4),

$$2\text{Im}(\mathcal{A}(i \rightarrow f)) = \sum_a (2\pi)^4 \delta^{(4)}(P_i - P_a) \mathcal{A}(f \rightarrow a)^* \mathcal{A}(i \rightarrow a). \quad (1.6.8)$$

This equality is called the optical theorem. Excluding overall factors of i , the only place that imaginary parts enter in the Feynman rules is in the $i\epsilon$ prescription for the Feynman propagator. Therefore, a nonzero $\text{Im}(\mathcal{A}(i \rightarrow f))$ can only arise when virtual particles go on-shell and the $i\epsilon$ becomes relevant.

The discontinuity of a function $f : \mathbb{C} \rightarrow \mathbb{C}$ at x_0 is given by the difference between the function evaluated at the point plus and minus a small imaginary part,

$$-i\text{Disc}_{x_0}(f(x)) := \lim_{\eta \rightarrow 0^+} \left(f(x_0 + i\eta) - f(x_0 - i\eta) \right), \quad (1.6.9)$$

where we include the $-i$ factor for later convenience. The discontinuity is only nonzero if it crosses a branch cut.

Consider our amplitude $\mathcal{A}(i \rightarrow f)$ as an analytic function of a complex variable, $\mathcal{A}(s)$. If the energy of state i in Eq. (1.6.8) is insufficient to produce on-shell intermediate states a , then $\text{Im}(\mathcal{A}(i \rightarrow f))$ vanishes. In this domain,

$$\mathcal{A}(s) = \mathcal{A}(s^*)^*. \quad (1.6.10)$$

Analytically continuing this to the entire complex s plane, in the domain where the initial energy is high enough to allow on-shell intermediate states and therefore $\text{Im}(\mathcal{A}(s))$ is nonzero,

$$\text{Re}(\mathcal{A}(s + i\epsilon)) = \text{Re}(\mathcal{A}(s - i\epsilon)), \quad (1.6.11)$$

$$\text{Im}(\mathcal{A}(s + i\epsilon)) = -\text{Im}(\mathcal{A}(s - i\epsilon)). \quad (1.6.12)$$

Thus, there is a branch cut across the real axis of s , showing that a nonzero $\text{Im}(\mathcal{A}(s))$ requires a branch cut singularity. The discontinuity is

$$\text{Disc}(\mathcal{A}(s)) := 2\text{Im}(\mathcal{A}(s)), \quad (1.6.13)$$

which can also be seen by applying Eqs. (1.6.6) and (1.6.10) to Eq. (1.6.9).

Examining Eq. (1.6.8) order by order in perturbation theory, we discover relations between the discontinuities of a given order and the products of lower-order terms

(dropping factors),

$$\text{Disc}\left(\mathcal{A}_n^{(0)}\right) = 0, \quad (1.6.14a)$$

$$\text{Disc}\left(\mathcal{A}_n^{(1)}\right) = \mathcal{A}_n^{(0)*} \mathcal{A}_n^{(0)}, \quad (1.6.14b)$$

$$\text{Disc}\left(\mathcal{A}_n^{(2)}\right) = \mathcal{A}_n^{(0)*} \mathcal{A}_n^{(1)} + \mathcal{A}_n^{(1)*} \mathcal{A}_n^{(0)} + \mathcal{A}_{n+1}^{(0)*} \mathcal{A}_{n+1}^{(0)}, \quad (1.6.14c)$$

and so on. The constraint Eq. (1.6.14b) is none other than the factorisation we found at one loop [147].

1.6.3 Generalised unitarity

Beyond those prescribed by unitarity, we can continue to make further cuts. This technique is called generalised unitarity [148, 149]. The full loop integrand can be reconstructed by analysing the complete set of possible unitarity cuts. This proceeds by identifying a basis of integrals, then isolating their coefficients by considering the projections provided by the cuts.

Consider n -point massless one-loop amplitudes with external states in four dimensions. They can be decomposed onto a basis of n -point scalar integrals $I_{k,i}^{(1)}$, labelled by the number of internal lines k , and a rational part R as

$$\mathcal{A}_n^{(1)} = \sum_{k=2}^4 \sum_i c_{k,i}^{(0)} I_{k,i}^{(1)} + R + \mathcal{O}(\epsilon). \quad (1.6.15)$$

The integrals are called boxes $I_{4,i}^{(1)}$, triangles $I_{3,i}^{(1)}$, and bubbles $I_{2,i}^{(1)}$. We neglect tadpoles $I_{1,i}^{(1)}$ since for massless propagators, the integrals are scaleless and hence vanish in DR. The size of the integral basis is set by the number of spacetime dimensions, four. This basis is presented in Ref. [150].

The coefficients $c_{k,i}^{(0)}$ are functions of the external kinematics and ϵ . Having distinct propagator structures, only a subset of the scalar integrals contribute to a given unitarity cut. Therefore, a single coefficient or set of coefficients may be singled out by performing a certain set of cuts. Quadruple cuts isolate only the box coefficients $c_{4,i}^{(0)}$ [151]; these are maximal cuts in $d = 4$. Triple cuts pick out both the box coefficients $c_{4,i}^{(0)}$ and the triangle coefficients $c_{3,i}^{(0)}$. Similarly, double cuts additionally include the bubble coefficients $c_{2,i}^{(0)}$. Thus, all coefficients can be extracted by solving a system of linear equations formed from various cuts. This gives the coefficients as combinations of products of trees [152–154].

Consider solution B. We can write helicity spinors which satisfy ℓ_1^B , ℓ_2^B , and ℓ_4^B through Eq. (1.5.40),

$$\begin{aligned} |\ell_1^B\rangle &= a_4|1\rangle, & |\ell_1^B] &= |2\rangle, \\ |\ell_2^B\rangle &= a_4|1\rangle - |2\rangle, & |\ell_2^B] &= |2\rangle, \\ |\ell_4^B\rangle &= |1\rangle, & |\ell_4^B] &= a_4|2\rangle + |1\rangle. \end{aligned} \tag{1.6.24}$$

We can treat solution A similarly. We now have all the ingredients we need to compute $c_{1^+|2^+|3^-|4^-5^-}^{(0)}$ as a function of the external kinematics.

Now we consider the trees of the cut one-loop diagram depicted in Eq. (1.6.17). Recall from Section 1.5.5 that three-gluon trees are nonzero for complex momenta. Since the trees are nonzero only for MHV helicity configurations, there is only a single free internal helicity, h_1 , with the others fixed,

$$h_2 = -, \quad h_3 = -, \quad h_4 = +. \tag{1.6.25}$$

Inserting our kinematic solutions into the MHV expressions for the trees Eq. (1.5.61), we find solution A vanishes, while solution B with $h_1 = -$ vanishes, leaving solution B with $h_1 = +$ as the only nonzero piece,

$$c_{1^+|2^+|3^-|4^-5^-}^{(0)} \Big|_{B, h_1=+} = \frac{[1 \ell_1^B]^3}{[\ell_1^B \ell_4^B][\ell_4^B 1]} \frac{\langle \ell_2^B \ell_1^B \rangle^3}{\langle \ell_1^B 2 \rangle \langle 2 \ell_2^B \rangle} \frac{\langle 3 \ell_3^B \rangle^3}{\langle \ell_3^B \ell_2^B \rangle \langle \ell_2^B 3 \rangle} \frac{\langle 4 5 \rangle^3}{\langle 5 \ell_4^B \rangle \langle \ell_4^B \ell_3^B \rangle \langle \ell_3^B 4 \rangle}. \tag{1.6.26}$$

The answer is given by the sum over internal helicities and average of the kinematic solutions, which after substitution of Eqs. (1.6.18) and (1.6.24) and subsequent algebra is

$$c_{1^+|2^+|3^-|4^-5^-}^{(0)} = \frac{1}{2} s_{12} s_{23} A^{(0)}(1_g^+, 2_g^+, 3_g^-, 4_g^-, 5_g^-). \tag{1.6.27}$$

The rational part R of Eq. (1.6.15) can be the most time-consuming to compute. While it is possible for the integral coefficients $c_{k,i}^{(0)}$, the rational part is not cut constructable with the loop-momentum in four dimensions. We can instead use cuts in analytic dimensions of $d = 4 - 2\epsilon$ [155–157]. As both the coefficients and integrals are Laurent series in ϵ , with the integrals containing poles, we see the origin of the rational part as terms in the expansion of the form

$$\mathcal{O}\left(\frac{1}{\epsilon^a}\right) \times \mathcal{O}(\epsilon^a), \quad a \in \mathbb{Z}^{\geq}. \tag{1.6.28}$$

Outside of four dimensions, the trees lose their simple structure. One way to manage this is to partition the four-dimensional part of the loop d -momentum, with the size of the extra dimensions introducing a new scale that we treat as a fictitious

mass,

$$k_{[d]}^\mu = k_{[4]}^\mu + k_{[-2\epsilon]}^\mu, \quad k_{[d]}^2 = k_{[4]}^2 - \mu^2, \quad \mu^2 = -k_{[-2\epsilon]}^2, \quad (1.6.29)$$

so that propagators look like massive four-dimensional propagators,

$$D_{\alpha_j}(k_{[d]}, p_i) = D_{\alpha_j}(k_{[4]}, p_i) - \mu^2. \quad (1.6.30)$$

1.6.4 Integrand reduction

We have discussed one-loop amplitudes in the form,

$$A_{n,j}^{(1)} = S_j a_j(p_i) \cdot \int \frac{d^d k}{(2\pi)^d} \frac{\mathcal{N}_j(k, p_i)}{\prod_{\alpha_j=1}^n D_{\alpha_j}(k, p_i)}, \quad (1.6.31)$$

with inverse propagators D_{α_j} defined in Eq. (1.6.5). We have factorised out from the numerator N_j the part depending only on external four-momenta a_j ,

$$N_j(k, p_i) = a_j(p_i) \cdot \mathcal{N}_j(k, p_i), \quad (1.6.32)$$

where \cdot represents the contraction of Lorentz indices. These are generally tensor integrals, as \mathcal{N}_j can carry free indices.

We can reduce to scalar integrals, meaning \mathcal{N}_j is scalar, by rewriting the numerator N_j in a basis of inverse propagators D_{α_j} and spurious irreducible scalar products (ISPs). An ISP is a scalar product that cannot be expressed in terms of only inverse propagators, while a spurious ISP vanishes after integration. A convenient choice is the van Neerven-Vermaseren basis [158].

For example, consider a massless box topology¹. By momentum conservation, there are only three LI external four-momenta. The four-dimensional components of the loop-momentum $k_{[4]}^\mu$ require one more orthogonal vector to find a basis. We introduce the vector,

$$\omega_{124}^\mu \propto \varepsilon^{\mu\nu\sigma\rho} (p_1)_\nu (p_2)_\sigma (p_4)_\rho, \quad (1.6.33)$$

such that it is orthogonal to an LI set of external four-momenta,

$$p_i \cdot \omega_{124} = 0 \quad \forall i \in \{1, 2, 4\}, \quad (1.6.34)$$

to form the basis,

$$k_{[4]}^\mu = \sum_i b_i v_i^\mu, \quad v = (p_1, p_2, p_4, \omega_{124}). \quad (1.6.35)$$

¹This example is inspired by the lecture series *Loop Amplitudes in Gauge Theories* given by Simon Badger at Durham University in 2019.

We can form the ISP,

$$k_{[4]} \cdot \omega_{124}, \quad (1.6.36)$$

which is spurious because

$$\int dk_{[4]} (k_{[4]} \cdot \omega_{124})^x = 0 \quad \forall x \in \mathbb{Z}^{\text{odd}}, \quad (1.6.37)$$

while other possible scalar products,

$$k_{[4]} \cdot p_i, \quad (1.6.38)$$

are reducible. For example,

$$\begin{aligned} k_{[4]} \cdot p_1 &= \frac{-(k_{[4]} - p_1)^2 + k_{[4]}^2}{2} \\ &= \frac{-D_2(k) + D_1(k)}{2}, \end{aligned} \quad (1.6.39)$$

so can be expressed as a linear combination of inverse propagators. We can then write an ansatz Δ_4 for a general box numerator. For a general renormalisable gauge theory, the maximum tensor rank for an n -propagator integral is n , so

$$\Delta_4 = y_0 + y_1 (k_{[4]} \cdot \omega_{124}) + y_2 \mu^2 + y_3 (k_{[4]} \cdot \omega_{124}) \mu^2 + y_4 \mu^4. \quad (1.6.40)$$

The coefficients y_1 and y_3 are spurious, y_0 is the four-dimensional scalar box coefficient, and y_2 and y_4 contain $\mathcal{O}(\epsilon)$ parts of the d -dimensional box cut. Choosing the index $j = 0$, this gives the integrand,

$$\partial \mathcal{A}_{4,0}^{(1)} = \frac{\Delta_4(k_{[4]} \cdot \omega_{124}, \mu^2)}{\prod_{\alpha_0=1}^4 D_{\alpha_0}} + \text{sub-topologies}. \quad (1.6.41)$$

After similarly treating the triangles and bubbles, we arrive at Eq. (1.6.15).

The connection to on-shell technology is provided by methods such as those from Ossola-Pittau-Papadopoulos (OPP) [159] or Forde [154], such that the integrand coefficients are projected out by numerical evaluations of cut diagrams.

1.6.5 Integration-by-parts identities

Having reduced a loop-level amplitude to scalar integrals, we may be interested in using a basis of LI scalar integrals, which we call master integrals (MIs) [160,161]. This reduces the number of integral evaluations required for the computation. Furthermore, the linear independence of the coefficients can be necessary for reconstruction techniques such as those discussed in Section 1.7. We can reduce to an MI basis by

using integration-by-parts (IBP) identities, also called recurrence relations, between the scalar integrals [160–162].

To illustrate IBP identities, consider a prototype one-loop scalar integral of the form,

$$I_a = \int \frac{d^d k}{(2\pi)^d} \frac{1}{(k^2 - m^2)^a}. \quad (1.6.42)$$

It is invariant under shifts of the loop-momentum k^μ , such as

$$k^\mu \rightarrow (1 + \zeta)k^\mu, \quad (1.6.43)$$

which gives the integral as

$$I_a = (1 + \zeta)^d \int \frac{d^d k}{(2\pi)^d} \frac{1}{([k(1 + \zeta)]^2 - m^2)^a}. \quad (1.6.44)$$

For infinitesimal shifts $\zeta \ll 1$, we can Taylor expand to find

$$I_a = I_a + \zeta \left(d I_a + \int \frac{d^d k}{(2\pi)^d} \frac{-2ak^2}{(k^2 - m^2)^{a+1}} \right) + \mathcal{O}(\zeta^2). \quad (1.6.45)$$

Using

$$\frac{\partial}{\partial k^\mu} k^\mu = d, \quad k^\mu \frac{\partial}{\partial k^\mu} \frac{1}{(k^2 - m^2)^a} = \frac{-2ak^2}{(k^2 - m^2)^{a+1}}, \quad (1.6.46)$$

we can rewrite Eq. (1.6.45) as

$$\frac{\partial}{\partial k^\mu} k^\mu I_a + \int \frac{d^d k}{(2\pi)^d} k^\mu \frac{\partial}{\partial k^\mu} \frac{1}{(k^2 - m^2)^a} = \mathcal{O}(\zeta^2). \quad (1.6.47)$$

Neglecting higher-order terms, we collect with the chain rule to obtain

$$\int \frac{d^d k}{(2\pi)^d} \frac{\partial}{\partial k^\mu} \left(\frac{k^\mu}{(k^2 - m^2)^a} \right) = 0. \quad (1.6.48)$$

This construction generalises to multi-loop integrals to give IBP identities of the form,

$$\int \left(\prod_\ell \frac{d^d k_\ell}{(2\pi)^d} \right) \frac{\partial}{\partial k_i^\mu} \left(\frac{v_i^\mu}{\prod_\alpha D_\alpha} \right) = 0, \quad (1.6.49)$$

where v_i^μ is some external or internal momentum and the D_α are the inverse propagators. The equation shows that the integral vanishes on the boundary, which it must for the result to be finite. This provides a method to generate relations between

different integrals.

A general method for identifying an MI basis and applying IBP identities to project onto it is the Laporta algorithm [163].

1.7 Finite fields

We now have the technology to decompose amplitudes into rational coefficients and special functions, and numerically evaluate the coefficients (Section 1.6). To obtain an optimal representation of the coefficients, we can reconstruct their analytic form through many numerical evaluations by randomly sampling an appropriate ansatz [164, 165]. Were we to perform this procedure using a floating-point number representation, the precision loss from rounding errors would spoil the computation. Instead, we sample over FFs, allowing us to use an integer number representation of fixed size. Unlike floats, integers have no precision loss¹ and we obtain an exact answer. We can extract the rational result from this computation at the end. This requires a rational parametrisation of the kinematics, which is provided by MTVs (Section 4.4). FF techniques have found much use in reconstruction and other methods in recent works [164–171].

1.7.1 Definition

An FF \mathbb{F}_n is a set of n non-negative integers,

$$\mathbb{F}_n = \{0, \dots, n-1\}, \quad (1.7.1)$$

together with the arithmetic operations of addition modulo n and multiplication modulo n , and their inverses. The size n of the field must be a prime power. The modular additive inverse is simply subtraction modulo n . The modular multiplicative inverse of a non-zero integer a is the integer x such that

$$ax = 1 \pmod{n}, \quad a \neq 0. \quad (1.7.2)$$

We use the following abuse of notation to denote the modular multiplicative inverse,

$$x = a^{-1} \pmod{n}. \quad (1.7.3)$$

This can be efficiently computed using the Extended Euclidean algorithm [172].

The modulus operation provides a non-injective homomorphism (a one-to-many

¹There is the caveat that the integer size is sufficiently large that overflows are not encountered. Working with operations modulo n for some integer n means that the calculation is safe if n is less than the integer size.

mapping) from the integers to an FF,

$$\text{mod } n : \mathbb{Z} \rightarrow \mathbb{F}_n . \tag{1.7.4}$$

Similarly, we can map from the rationals $\mathbb{Q} \rightarrow \mathbb{F}_n$ with

$$\frac{a}{b} \rightarrow a b^{-1} \text{ mod } n . \tag{1.7.5}$$

1.7.2 Reconstruction

Let us demonstrate the idea of FF reconstruction with an example¹. Consider a box contribution to the one-loop six-gluon helicity amplitude,

$$\mathcal{A}_6^{(1)}(1^+, 2^-, 3^+, 4^-, 5^+, 6^-) = d^{1^+|2^-|3^+4^-5^+|6^-} \left(\begin{array}{c} \text{Diagram: A square loop with external lines 1, 2, 3, 4, 5, 6. Line 1 is top-left, 2 is top-right, 3 is right, 4 is bottom-right, 5 is bottom, 6 is left.} \end{array} \right) + \dots , \tag{1.7.6}$$

where

$$d^{1^+|2^-|3^+4^-5^+|6^-} = \sum_{h_1, h_2, h_3, h_4 \in \{+, -\}} d_{h_1 h_2 h_3 h_4}^{1^+|2^-|3^+4^-5^+|6^-} . \tag{1.7.7}$$

The lines in the superscript of the coefficient indicate the cuts, while the subscript shows the helicities in the loop, as shown in the cut diagram,

$$d_{h_1 h_2 h_3 h_4}^{1^+|2^-|3^+4^-5^+|6^-} = \left(\begin{array}{c} \text{Diagram: A square loop with external lines 1^+, 2^-, 3^+, 4^-, 5^+, 6^-. Internal lines are labeled h1, h2, h3, h4 with arrows indicating helicity flow.} \end{array} \right) . \tag{1.7.8}$$

In particular, consider the coefficient for $\{h_1, h_2, h_3, h_4\} = \{+, -, +, -\}$, which takes the form,

$$d_{+-+-}^{1^+|2^-|3^+4^-5^+|6^-} = \frac{\mathcal{N}}{[34][45]\langle 2|(1+6)|5\rangle\langle 6|(1+2)|3\rangle s_{345}} . \tag{1.7.9}$$

¹This example is inspired by a question set by Ben Page at the SAGEX Mathematica and Maple School 2021.

The numerator \mathcal{N} can be written in terms of the spinor brackets, so we can write the ansatz,

$$\mathcal{N} = \sum_i \alpha_i g_i, \quad g_i = \prod_{j=1}^{30} v_j^{\beta_{ij}}, \quad \beta_{ij} \in \mathbb{Z}^>, \quad (1.7.10)$$

such that the numerator is a linear combination of polynomials g_i in the spinor brackets v_j ,

$$\vec{v} = \{ \langle ij \rangle, [ij] : i < j; i, j \in \{1, \dots, 6\} \}. \quad (1.7.11)$$

By Eq. (1.3.7) and Eq. (1.5.58), we know the mass dimension and little group scalings of the numerator,

$$[\mathcal{N}] = 10, \quad \{\mathcal{N}\} = \{-2, 3, -4, 0, -4, 3\}. \quad (1.7.12)$$

This provides physical constraints on the ansatz,

$$\sum_{j=1}^{30} \beta_{ij} = [\mathcal{N}], \quad \sum_{j=1}^{30} \beta_{ij} \{v_j\}_k = \{\mathcal{N}\}_k \quad \forall i, k, \quad (1.7.13)$$

since

$$[v_j] = 1, \quad [\alpha_i] = 0, \quad [g_i] = [\mathcal{N}]. \quad (1.7.14)$$

There are 93 solutions to the linear system of constraints Eq. (1.7.13), so $i \in \{1, \dots, 93\}$. However, the solutions are not all LI. To reduce to the LI subset, we numerically sample 93 phase-space points $\{x_i\}_{i=1}^{93}$ to construct the 93×93 matrix,

$$M_1 = \begin{pmatrix} g_1(x_1) & \dots & g_{93}(x_1) \\ \vdots & \ddots & \vdots \\ g_1(x_{93}) & \dots & g_{93}(x_{93}) \end{pmatrix}. \quad (1.7.15)$$

We randomly generate the points x_i using a momentum twistor parametrisation such that they are rational and map them to a finite field $\mathbb{Q} \rightarrow \mathbb{F}_{n_1}$. We perform row reduction, also called Gaussian elimination, (modulo n_1) on the matrix M_1 to identify the 24 LI polynomials \tilde{g}_i by the position of the leading number one in the upper 24 rows of the matrix in reduced row echelon form. This reduces the ansatz to

$$\mathcal{N}(x) = \sum_{i=1}^{24} a_i \tilde{g}_i(x). \quad (1.7.16)$$

To solve for the LI coefficients a_i , we numerically sample the components of

$$0 = -\mathcal{N}(x) + \sum_{i=1}^{24} a_i \tilde{g}_i(x) \quad (1.7.17)$$

over 25 points $\{y_i\}_{i=1}^{25}$ to form the 25×25 matrix,

$$M_2 = \begin{pmatrix} \mathcal{N}(y_1) & \tilde{g}_1(y_1) & \dots & \tilde{g}_{25}(y_1) \\ \vdots & \vdots & \ddots & \vdots \\ \mathcal{N}(y_{25}) & \tilde{g}_1(y_{25}) & \dots & \tilde{g}_{25}(y_{25}) \end{pmatrix}, \quad (1.7.18)$$

and find the nullspace, also called the kernel, (modulo n_1) of M_2 . In other words, we solve

$$M_2 \vec{a} = \vec{0} \quad (1.7.19)$$

for the vector $\vec{a} \in (\mathbb{F}_{n_1})^{25}$. After multiplying modulo n_1 the components \tilde{a}_i by

$$-\tilde{a}_0^{-1} \pmod{n_1} \quad (1.7.20)$$

to fix the coefficient of \mathcal{N} as -1 , we find the form

$$\vec{a} = \{-1, a_1, \dots, a_{24}\}, \quad (1.7.21)$$

with numerical values for the a_i in \mathbb{F}_{n_1} . Thus, having the values of the coefficients of our ansatz Eq. (1.7.16), we have found \mathcal{N} in \mathbb{F}_{n_1} .

1.7.3 Recovering rationals

We now want to recover the rational solution for \mathcal{N} , but the map $\mathbb{F}_{n_1} \rightarrow \mathbb{Q}$ is many-to-one. For each coefficient a_i , we could solve

$$\frac{r_i}{s_i} \pmod{n_1} = a_i \quad (1.7.22)$$

and fix the ambiguity by choosing the rational r_i/s_i with smallest Euclidean norm $r_i^2 + s_i^2$. However, this introduces a lower bound on n_1 as it requires that

$$r_i^2 + s_i^2 < n_1 \quad (1.7.23)$$

and we don't know how large the norm will be a priori. In addition, if n_1 were required to be too large, we would lose computational efficiency by being forced to use non-native integer number representation.

The solution is provided by the Chinese Remainder Theorem, which allows us to construct a result in a larger FF from multiple evaluations in smaller FFs. Let us

demonstrate this by applying it to our reconstruction example from Section 1.7.2, labelling the components of the result Eq. (1.7.21) as $a_i^{(n_1)}$. We repeat the steps in Section 1.7.2 using a different FF \mathbb{F}_{n_2} , where n_1 and n_2 are coprime, to obtain the $a_i^{(n_2)}$. These numbers are related through a number b_i , in our case the rational result, by

$$a_i^{(n_1)} = b_i \pmod{n_1} \quad a_i^{(n_2)} = b_i \pmod{n_2} \quad (1.7.24a)$$

$$\implies b_i = a_i^{(n_1)} + m_{1,i}n_1 \quad b_i = a_i^{(n_2)} + m_{2,i}n_2, \quad (1.7.24b)$$

for some $m_{1,i}, m_{2,i} \in \mathbb{Z}^{\geq}$. Since n_1 and n_2 are coprime, Bézout's identity tells us that there exist $q_1, q_2 \in \mathbb{Z}$ such that

$$1 = q_1n_1 + q_2n_2. \quad (1.7.25)$$

Multiplying through by b_i and inserting from Eq. (1.7.24b),

$$\begin{aligned} b_i &= \left(a_i^{(n_2)} + m_{2,i}n_2\right) q_1n_1 + \left(a_i^{(n_1)} + m_{1,i}n_1\right) q_2n_2 \\ &= \underbrace{a_i^{(n_1)}q_2n_2 + a_i^{(n_2)}q_1n_1}_{a_i^{(n_1n_2)}} + \underbrace{\left(q_1m_{2,i} + q_2m_{1,i}\right)}_{m_{3,i}} n_1n_2. \end{aligned} \quad (1.7.26)$$

Hence, we have constructed the result in $\mathbb{F}_{n_1n_2}$,

$$a_i^{(n_1n_2)} = b_i \pmod{(n_1n_2)}. \quad (1.7.27)$$

This method generalises to evaluations over an arbitrary number of FFs. Choosing the FF sizes n_j as prime numbers ensures all n_j are coprime. This provides a way to construct a result $a_i^{(\prod_j n_j)}$ in a sufficiently large FF $\mathbb{F}_{\prod_j n_j}$ that

$$r_i^2 + s_i^2 < \prod_j n_j \quad (1.7.28)$$

is satisfied for the recovery of the rational result r_i/s_i , while keeping evaluations over sufficiently small FFs \mathbb{F}_{n_j} for computational efficiency.

1.8 Phenomenology

Precise theoretical predictions are in high demand for the current LHC experiments, which are aiming to look for tiny deviations from the SM. These include the *A Toroidal LHC Apparatus* (ATLAS) [173] and *Compact Muon Solenoid* (CMS) [174] experiments. With experimental bottlenecks like the determination of interaction luminosity at around one percent at ATLAS [175] and CMS [176], and similarly for the resolution of jet energies [177,178], the current target for theory is to also achieve

one percent precision [179]. In the near future, the High Luminosity upgrade of the LHC will also overcome statistical limitations [180].

Due to the relatively large size of the strong coupling constant, NNLO corrections in QCD are desirable for a wide variety of final state processes. In particular, a class of two-to-three scattering processes with many kinematic scales have presented a considerable challenge to the theoretical community and there has been a good deal of activity leading to new methods able of overcoming their algebraic and analytic complexity [131, 171, 181–189]. We discuss recent progress for diphoton-plus-jet production in Section 1.8.3.

1.8.1 Fixed-order perturbation theory

In fixed-order perturbation theory, in analogy to our expansion of the amplitude, Eq. (1.3.5), we calculate the cross section as an expansion in the coupling. For example, up to NNLO,

$$d\sigma = \alpha_s^A (d\sigma_{\text{LO}} + \alpha_s d\sigma_{\text{NLO}} + \alpha_s^2 d\sigma_{\text{NNLO}}) + \mathcal{O}(\alpha_s^{A+3}), \quad (1.8.1)$$

where A is the coupling power of the LO term. It is important to calculate order-by-order for the regularisation of IR divergences, as we will see in Section 2.3. Since $\alpha_s \approx 0.1$ at the energy scales of modern colliders, we can estimate that at least NNLO predictions are necessary in general for percent-level precision.

To evince the anatomy of this series, let us consider the ME contributions (up to permutations and antiparticles) to hadronic diphoton-plus-jet production $pp \rightarrow \gamma\gamma j$:

LO At LO, which is also called Born level, we have $A = 1$. The contributions are the square of the amplitudes $q\bar{q} \rightarrow \gamma\gamma g$ and $qg \rightarrow \gamma\gamma q$. We do not have the gluon-fusion channel as there is no vertex coupling photons to gluons in the SM.

NLO Moving to NLO, we can increase the order in α_s by manipulating the LO diagrams in two ways: either by emitting an additional particle in the final state, called a real correction; or by introducing a new internal line to form a closed loop, called a virtual correction. Note that both of these corrections are experimentally indistinguishable from the Born as virtual particles are not observed and the final state reconstruction is dependent on the jet clustering algorithm. The real correction includes the square of the amplitudes $q\bar{q} \rightarrow \gamma\gamma gg$, $qg \rightarrow \gamma\gamma qg$, and $gg \rightarrow \gamma\gamma q\bar{q}$, while the virtual correction is the interference of the tree-level and one-loop diagrams for $q\bar{q} \rightarrow \gamma\gamma g$ and $qg \rightarrow \gamma\gamma q$. Note that the virtual contribution to the cross section involves an integration over




Figure 1.6: Schematic of the perturbative expansion in α_s of $pp \rightarrow \gamma\gamma j$. Blobs represent sums over possible internal vertex configurations for the external particles, with holes denoting loops. Single representative terms are shown to highlight each contribution, with other terms omitted. Moving to the right horizontally, we have an increase in the number of loops (virtual corrections). Moving downwards vertically, we have an increase in the number of additional emissions (real corrections). LO (green) is $\mathcal{O}(\alpha_s)$. At NLO (yellow), which is $\mathcal{O}(\alpha_s^2)$, there are real and virtual contributions. At NNLO (orange), which is $\mathcal{O}(\alpha_s^3)$, there are double-real (RR), real-virtual (RV), and VV contributions. In the VV, we see the first appearance of the gluon-fusion subprocess (blue). At (next-to-)³leading order (N³LO) (red), which is $\mathcal{O}(\alpha_s^4)$, there are triple-real (R³), real-real-virtual (RRV), real-virtual-virtual (RVV), and triple-virtual (V³) contributions. The NLO contributions to the gluon-fusion subprocess are highlighted here in cyan.




Figure 1.7: Plot of the (k_F, k_R) scale factor set for the seven-point scale variation.

a three-particle phase space, while the real has an extra particle and thus is over a four-particle phase space.

N^kLO We can continue this pattern to higher orders, (next-to-)^kleading order (N^kLO), as depicted in Fig. 1.6. The first place that the gluon-fusion subprocess $gg \rightarrow \gamma\gamma g$ can appear is in the VV correction at NNLO of the full process $pp \rightarrow \gamma\gamma j$. Because it has no tree-level diagram and thus contains a loop in the LO contribution, we refer to $gg \rightarrow \gamma\gamma g$ as a loop-induced process.

1.8.2 Estimating uncertainty

Theory uncertainty

Fixed-order theoretical predictions come with an error due to the missing higher orders. While the only way to really know this uncertainty is to calculate the next order, in practice we can estimate it. As the scale dependence is due to missing higher orders (Section 1.3.2), and a scale variation amounts to an $\mathcal{O}(\alpha_s)$ correction, we use it as a measure of the uncertainty due to the missing higher orders.

The de facto choice is an ad hoc method known as the seven-point scale variation. The computation for an observable $\sigma(\mu_F, \mu_R)$ is repeated for a set of different scales varied around a central value. Typically, the central value is taken as

$$\mu_F = \mu_R =: \mu, \quad (1.8.2)$$

with μ set to a fixed scale or calculated dynamically as in Eq. (5.3.2). We consider a set of values of (k_F, k_R) in

$$S_7 = \{(0.5, 0.5), (0.5, 1), (1, 0.5), (1, 1), (1, 2), (2, 1), (2, 2)\}, \quad (1.8.3)$$

such that we sample the seven cases of $\sigma(k_F\mu, k_R\mu)$ where

$$k_F, k_R \in \{0.5, 1, 2\}, \quad k_F/k_R \in \{0.5, 1, 2\}; \quad (1.8.4)$$

this is depicted in Fig. 1.7. The result is then that of the central value $\sigma(\mu, \mu)$, with the error estimated by the interval

$$\left[\min\left(\sigma\left(k_F\mu, k_R\mu\right) : (k_F, k_R) \in S_7\right), \max\left(\sigma\left(k_F\mu, k_R\mu\right) : (k_F, k_R) \in S_7\right) \right]. \quad (1.8.5)$$

There are also efforts to prescribe more reliable methods of uncertainty estimation, such as Bayesian modelling [190–192].

Monte Carlo uncertainty

The phase-space integration involved in the computation of an observable like the cross section in a collider simulation is generally performed using the numerical method of MC integration [193]. This technique approximates an integral over a function $f(\vec{x})$,

$$I = \int_V d^D x f(\vec{x}), \quad (1.8.6)$$

as the sum,

$$I \approx \frac{V}{N} \sum_{i=1}^N f(\vec{x}_i), \quad (1.8.7)$$

where D is an integer and the D -dimensional phase space of finite volume V is sampled N times at points \vec{x}_i with a uniform random distribution. The error scales like $1/\sqrt{N}$ and can be estimated by the sample standard error [194],

$$\alpha = \frac{\sigma_{N-1}}{\sqrt{N}}, \quad (1.8.8)$$

which we define through the sample standard deviation,

$$\sigma_{N-1} = \sqrt{\frac{1}{N-1} \sum_{i=1}^N (f(x_i) - \bar{f})^2}, \quad (1.8.9)$$

and the mean,

$$\bar{f} = \frac{1}{N} \sum_{i=1}^N f(x_i). \quad (1.8.10)$$

By the Central Limit Theorem [194], the MC result converges on the true result as $N \rightarrow \infty$. Various methods such as importance sampling and stratified sampling are used to accelerate convergence.

For differential observables, the MC error on the i^{th} histogram bin, which contains points in the set b_i , can be estimated by

$$\alpha_i = \sqrt{\frac{1}{N(N-1)} \sum_{j=1}^N (\theta(j \in b_i) f(x_j) - \bar{f}_i)^2}, \quad (1.8.11)$$

with

$$\bar{f}_i = \frac{1}{N} \sum_{j=1}^N \theta(j \in b_i) f(x_j), \quad (1.8.12)$$

where N is the total number of points in all bins. To ensure the area under the histogram is equal to the total integrated observable, the bin heights and errors can be divided by the bin width.

Histograms can be rebinned by aggregating adjacent bins. Assuming the samples to be random and uncorrelated, the original bin errors can be propagated to the new bins by adding errors in quadrature. The resultant bin error is a lower bound. To propagate errors exactly in rebinning, the covariance matrix C can be used,

$$C_{ik} = \frac{1}{N-1} \sum_{j=1}^N (\theta(j \in b_i) f(x_j) - \bar{f}_i) (\theta(j \in b_k) f(x_j) - \bar{f}_k) \quad (1.8.13)$$

where the indices i and k denote the bins. The diagonal elements of the covariance matrix give the variances and are related to the bin errors by

$$\alpha_i = \sqrt{\frac{C_{ii}}{N}}, \quad (1.8.14)$$

since the variance is the square of the standard deviation, while the off-diagonal terms give the correlations between bins. When aggregating bins, for example from four to two bins by combining the first and second pair, the covariance matrix is aggregated as

$$\left(\begin{array}{cc|cc} c_{11} & c_{12} & c_{13} & c_{14} \\ c_{21} & c_{22} & c_{23} & c_{24} \\ \hline c_{31} & c_{32} & c_{33} & c_{34} \\ c_{41} & c_{42} & c_{43} & c_{44} \end{array} \right) \rightarrow \left(\begin{array}{cc} c_{11} + c_{12} + c_{21} + c_{22} & c_{13} + c_{14} + c_{23} + c_{24} \\ c_{31} + c_{32} + c_{41} + c_{42} & c_{33} + c_{34} + c_{43} + c_{44} \end{array} \right). \quad (1.8.15)$$

The variances for the new bins are given by the diagonal entries.

1.8.3 Diphoton production

The production of a pair of high-energy photons is an important experimental signature at hadron colliders and can be used, for example, to study the Higgs boson through its decay to photons [195] or in beyond the SM searches [196]. The SM background is dominated by QCD corrections, and a precise description of the kinematics of these observables requires the theoretical predictions to include perturbative information from the production in association with additional jets. Diphoton-plus-jet signatures form the largest background to Higgs production at high transverse momenta; the extra jet is necessary to ensure a non-zero transverse momentum in the diphoton system. In addition, with at least one jet, the Higgs coupling can be probed independently of initial-state effects, thus bypassing additional uncertainties introduced by the PDFs.

Diphoton production $pp \rightarrow \gamma\gamma$ has been known at NNLO for some time [197–199] and the two-loop scattering amplitudes were among the first complete $2 \rightarrow 2$ processes to be calculated at this order [200, 201]. A q_T -resummed calculation at order¹ NNLO + N³LL' was presented recently [202]. Steps towards N³LO are being taken with the completion of the first four-point three-loop QCD amplitude [203]. For the gluon-fusion subprocess $gg \rightarrow \gamma\gamma$, the three-loop amplitudes were recently calculated [204]. Along with the five-point two-loop results of Chapter 4, this makes available the final remaining piece for the NNLO corrections to the subprocess, which contribute to the full process at (next-to-)⁴leading order (N⁴LO).

NNLO QCD corrections to diphoton-plus-jet production $pp \rightarrow \gamma\gamma j$, which is initiated at LO by quark-antiquark and quark-gluon processes (Section 1.8.1), have been considered a high priority for current and future experiments for several years [205–207]. Following recent breakthroughs in two-loop amplitude technology, there has been a flurry of activity around this channel, including complete sets of LC [208, 209] and now full-colour (FC) [210] two-loop helicity amplitudes, and LC NNLO distributions [211]. This progress is extremely timely given the continually improving experimental measurements of diphoton signatures [212].

The loop-induced gluon-fusion channel $gg \rightarrow \gamma\gamma g$ starts contributing to the cross section only from NNLO onwards. Owing to the large gluon luminosity, *i.e.* the gluon PDF diverges at small momentum fraction x in high-energy scales (see Fig. 1.3), this channel yields a dominant contribution to the NNLO corrections and dominates their scale uncertainty [211]. To improve upon this uncertainty requires the NLO corrections, for which we derive the FC two-loop virtual amplitudes in Chapter 4 and the FC NLO distributions in Chapter 5.

¹A prime indicates the inclusion of partial but dominant corrections from the next higher order.

1.9 Structure of this thesis

In Chapter 2, we study the IR factorisation properties of QCD amplitudes and discuss how IR divergences are regulated in fixed-order cross section calculations before presenting a library of IR singular functions. In Chapter 3, we review the use of neural networks (NNs) to emulate MEs for processes with two photons and many gluons in hadron collider simulations. In Chapter 4, we present the first full set of FC five-point helicity amplitudes at two loops in QCD, which are those for the process of diphoton-plus-jet production through gluon fusion. In Chapter 5, we combine these amplitudes with the real contributions to produce NLO QCD distributions, the first time that FC two-to-three two-loop amplitudes have been integrated to provide fully differential cross section predictions relevant for the LHC. In Chapter 6, we also present a set of LC two-loop helicity amplitudes for trijet production at hadron colliders. We conclude in Chapter 7.

Chapter 2

Infrared behaviour

In this chapter, we look into the IR behaviour of QCD amplitudes and cross sections. In Section 2.1, we discuss the soft and collinear limits that can appear in phase space, and in Section 2.2, the factorisation properties of amplitudes in these limits. We motivate the utility of these limits in Section 2.2.4, and explain how they are regulated in the calculation of observables in Section 2.3, also listing proposed schemes for NNLO. In Section 2.4, we present a library of universal IR singular helicity functions. For further reading, see the review [213].

2.1 Infrared limits

We saw in Section 1.3.2 that loops contain IR divergences when the loop-momentum becomes small, which we regulate with DR. IR divergences also manifest in amplitudes when the phase space of the external momenta contains unresolved particles, which occurs in soft and collinear limits. In the following, we take the massless four-momentum as $p = (E, \vec{p})$. The double-collinear limit, denoted $i \parallel j$, occurs when $\theta_{ij} \rightarrow 0$ where θ_{ij} is the angle between \vec{p}_i and \vec{p}_j . Similarly, the single-soft limit \mathcal{E}_i of a massless particle i is when $E_i \rightarrow 0$, or equivalently when $|\vec{p}_i| \rightarrow 0$. These configurations cause an internal propagator to go on-shell, giving rise to a singularity.

The kinematics of the collinear limit $i \parallel j$ are characterised by $s_{ij} \rightarrow 0$ and the soft limit \mathcal{E}_i by $s_{ij} \rightarrow 0 \forall j$ since

$$s_{ij} = 2|\vec{p}_i||\vec{p}_j|(1 - \cos(\theta_{ij})) . \quad (2.1.1)$$

Soft limits and collinear limits can overlap in phase space, so care must be taken to avoid double counting when considering them.

2.1.1 Cuts

At LO, there exist IR phase-space regions. However, these are not physical states as they cannot be measured: soft particles are invisible and collinear particles are indistinguishable. Thinking of detectors, there is a finite resolution to which particles can be measured; in theory, our phase space should be similarly constrained. This is achieved by setting appropriate cuts, which are criteria that phase-space points must satisfy.

Some quantities to cut on, which we will make use of later in this thesis, include the following.

- Transverse momentum magnitude,

$$p_T = \sqrt{p_x^2 + p_y^2}, \quad (2.1.2)$$

with the beam along the z -axis.

- Pseudorapidity,

$$\eta = -\ln\left(\tan\left(\frac{\theta}{2}\right)\right), \quad (2.1.3)$$

where the polar angle θ is between the three-momentum of the particle \vec{p} and the positive direction of the beam axis \hat{z} .

- Azimuthal angle ϕ .

Related to cuts is the concept of jet clustering. A common choice for identifying QCD jets is the anti- k_T algorithm [214], which is implemented in `FastJet` [215]. Considering QED states, photons γ are often selected in an IR safe fashion by Frixione smooth cone isolation [216]. It prescribes that all cones of radius $r_\gamma < \Delta R$ must satisfy

$$E_{\text{hadronic}}(r_\gamma) \leq \epsilon_\gamma E_{T,\gamma} \frac{1 - \cos(r_\gamma)}{1 - \cos(\Delta R)},$$

where $E_{\text{hadronic}}(r_\gamma)$ is the hadronic energy found inside the cone of radius r_γ , $E_{T,\gamma}$ is the transverse momentum of the photon, the isolation cut cone radius is given by the R-separation,

$$\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}, \quad (2.1.4)$$

and ϵ_γ is an arbitrary parameter.

2.1.2 Poles

Beyond LO, we calculate contributions to the fixed-order result that are not themselves observable and thus really contain IR divergences. For instance, at NLO, the real correction contains single-unresolved singularities, and the loop in the amplitude for the virtual correction diverges when the loop-momentum goes soft. In DR, the one-loop amplitude therefore contains poles in the dimensional regulator ϵ ; when expanded as a Laurent series in ϵ , it contains coefficients of ϵ^{-2} and ϵ^{-1} , and a finite part, with $\mathcal{O}(\epsilon)$ terms being neglected¹ as they vanish in $\epsilon \rightarrow 0$. However, the poles in the real and virtual corrections come with opposite signs and exactly cancel. In fact, for any IR safe observable, when all components at any fixed order are combined, the result is finite, as described by the Kinoshita-Lee-Nauenberg (KLN) theorem [217, 218].

It can also be the case, at any order, that certain terms within an amplitude diverge in the IR limit of some kinematic variable or combination thereof, but do not correspond to a physical IR singularity. These crop up depending on the particular algebraic expression of the amplitude and cancel out in the full finite result. As they are unphysical, they are called spurious poles. Performing partial fraction decomposition (PFD) can introduce spurious poles; it can be necessary to take care with them as they can introduce large intermediate cancellations that reduce numerical stability. An example is considered in Section 4.5.3.

2.2 Infrared factorisation

An all-orders $(n + m)$ -particle amplitude factorises in the m -unresolved IR limit as

$$\mathcal{A}_{n+m} \rightarrow \tilde{\mathcal{A}}_n \otimes \mathcal{S}_m, \quad (2.2.1)$$

where: $\tilde{\mathcal{A}}_n$ is the reduced amplitude, which is none other than the on-shell n -particle amplitude; \otimes indicates summation over spin and colour states of correlated particles; and \mathcal{S}_m is a universal singular function for the m -unresolved limit. The double-collinear and single-soft limits described in Section 2.1 are single-unresolved limits. If there is no summation over correlated states, *i.e.* the \otimes reduces to an ordinary multiplication \times , this is called an exact factorisation. For QCD amplitudes, it is the partial amplitudes that exhibit this factorisation property, for which only adjacent legs can become singular.

¹Note that we need up to $\mathcal{O}(\epsilon^2)$ for the amplitudes in the one-loop squared contributions to VV corrections at NNLO as crossterms in the square yield poles and finite parts.

2.2.1 Collinear limits

Consider the tree-level single-unresolved IR limit when two partons go collinear. For an $(n + 1)$ -parton amplitude, there is a factorisation into an n -parton reduced amplitude and a three-legged splitting amplitude,

$$\mathcal{A}_{n+1}^{(0)} \rightarrow \tilde{\mathcal{A}}_n^{(0)} \otimes \mathcal{P}_1^{(0)}. \quad (2.2.2)$$

The $i \parallel j$ phase-space configuration can be constructed by the parametrisation,

$$p_i^\mu = z \tilde{p}^\mu + p_\perp^\mu - \frac{p_\perp^2}{z 2 \tilde{p} \cdot q} q^\mu, \quad (2.2.3a)$$

$$p_j^\mu = (1 - z) \tilde{p}^\mu - p_\perp^\mu - \frac{p_\perp^2}{(1 - z) 2 \tilde{p} \cdot q} q^\mu, \quad (2.2.3b)$$

$$\tilde{p}^\mu = p_{ij}^\mu - \frac{s_{ij}}{2 p_{ij} \cdot q} q^\mu, \quad (2.2.3c)$$

$$z = \frac{p_i \cdot q}{p_{ij} \cdot q}, \quad s_{ij} = -\frac{p_\perp^2}{z(1 - z)}, \quad (2.2.3d)$$

where: the transverse momentum p_\perp is timelike ($p_\perp^2 < 0$) and orthogonal to q , \tilde{p} , and p_{ij} ; \tilde{p} is the massless projection of p_{ij} ; and q is a reference null (light-like) four-vector ($q^2 = 0$). The parameter z controls the momentum fraction of the unresolved pair in leg i , and $|p_\perp| \rightarrow 0$ gives the collinear limit. We can write this as

$$\{\dots, p_i, p_j, \dots, p_k, \dots\}_{n+1} \xrightarrow{i \parallel j} \{\dots, \tilde{p}, \dots, \tilde{p}_k, \dots\}_n, \quad (2.2.4)$$

where p_i and p_j live in the full $(n + 1)$ -particle phase space, and map to the on-shell momentum p in the reduced n -particle phase space. When not in the exact limit, the reduced phase space must be provided through a mapping of the full phase space that imposes on-shellness and momentum conservation. We also select an arbitrary momentum p_k from the full phase space that we recoil off in the mapping to obtain \tilde{p}_k in the reduced phase space. One such mapping is described in the Catani-Seymour (CS) dipole subtraction scheme (Section 2.3), in which they call the correlated leg \tilde{p} the emitter and the recoiling leg \tilde{p}_k the spectator. It can be stated as

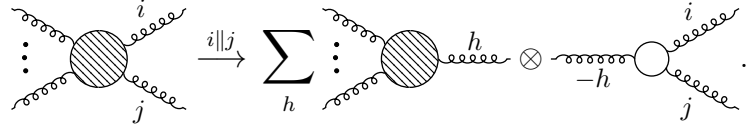
$$\tilde{p}_k^\mu = \left(1 + \frac{s_{ij}}{2 p_{ij} \cdot q}\right) p_k^\mu, \quad q = p_k. \quad (2.2.5)$$

Another mapping is provided by Kosower's antenna scheme [219].

A helicity amplitude in the $i \parallel j$ limit factorises as

$$\mathcal{A}_n^{c_1 \dots c_{i-1} c_i c_j c_{j+1} \dots c_n}(1^{h_1}, \dots, (i-1)^{h_{i-1}}, i^{h_i}, j^{h_j}, (j+1)^{h_{j+1}}, \dots, n^{h_n}) \xrightarrow{i \parallel j} \sum_{h_\rho, c_\rho} \tilde{\mathcal{A}}_{n-1}^{c_1 \dots c_{i-1} c_\rho c_{j+1} \dots c_n}(1^{h_1}, \dots, (i-1)^{h_{i-1}}, \rho^{h_\rho}, (j+1)^{h_{j+1}}, \dots, n^{h_n}) \mathcal{P}_1^{c_i c_\rho c_j}(i^{h_i}, \rho^{-h_\rho}, j^{h_j}), \quad (2.2.6)$$

where c_i is the abstract colour index (*i.e.* we have not specified the representation) associated with leg i , $h_i \in \{-, +\}$ is the helicity of leg i , and ρ labels the correlated leg. There remains a sum over the helicity h_ρ and the colour c_ρ . Pictorially, drawing gluons in lieu of any parton and with colour implicit, this looks like



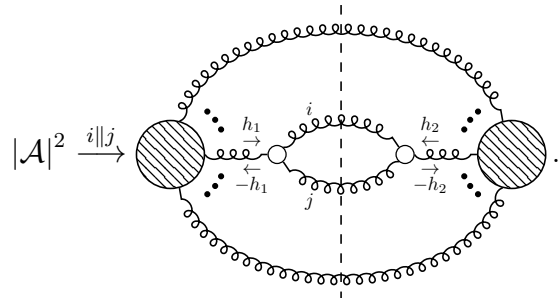
$$\xrightarrow{i \parallel j} \sum_h \dots \otimes \dots \quad (2.2.7)$$

With this picture in mind, we see that the double-collinear partonic splitting amplitudes, up to antiparticles and permutations, are $g \rightarrow gg$ and $q \rightarrow qg$ (giving $g \rightarrow \bar{q}q$ by crossing symmetry). The splitting amplitudes can be algebraically derived by constructing these configuration with MHV amplitudes [220], for example.

Turning to the squared helicity amplitude, we have

$$|\mathcal{A}|^2 \xrightarrow{i \parallel j} \left(\sum_{h_1, c_\rho} \tilde{\mathcal{A}}_{h_1}^{c_1 \dots c_\rho \dots c_{n_g}} \mathcal{P}_{-h_1}^{c_i c_\rho c_j} \right)^\dagger \sum_{h_2, c_\sigma} \tilde{\mathcal{A}}_{h_2}^{c_1 \dots c_\sigma \dots c_{n_g}} \mathcal{P}_{-h_2}^{c_i c_\sigma c_j}, \quad (2.2.8)$$

where subscripts now denote the helicity of the correlated leg. We can represent this diagrammatically as



$$|\mathcal{A}|^2 \xrightarrow{i \parallel j} \dots \quad (2.2.9)$$

Consider the colour of the squared splitting amplitude,

$$\left(\mathcal{P}_{-h_1}^{a_i a_\rho a_j} \right)^\dagger \mathcal{P}_{-h_2}^{a_i a_\sigma a_j} \Big|_{\text{colour}} = \begin{cases} f^{a_i a_\rho a_j} f^{a_i a_\sigma a_j} = C_A \delta^{a_\rho a_\sigma} & \text{for } g \rightarrow gg, \\ t_{a_i a_j}^{a_\rho} t_{a_i a_j}^{a_\sigma} = T_F \delta^{a_\rho a_\sigma} & \text{for } g \rightarrow \bar{q}q, \\ t_{a_i a_\rho}^{a_j} t_{a_i a_\sigma}^{a_j} = C_F \delta_{a_\rho a_\sigma} & \text{for } q \rightarrow qg. \end{cases} \quad (2.2.10)$$

Thus, contracting the colour indices of the splitting amplitudes always results in a Kronecker delta of the collinear colour indices (Section 1.4.1). This breaks the colour correlation that we observed at amplitude level to give exact factorisation of colour in the squared amplitude. Armed with this observation, Eq. (2.2.8) can be rearranged as

$$|\mathcal{A}|^2 \xrightarrow{i||j} \sum_{h_1, h_2} \widetilde{\mathcal{M}}_{h_1 h_2} \mathcal{P}_{-h_1 - h_2}, \quad (2.2.11)$$

where we have introduced the 2×2 spin matrices,

$$\mathcal{P}_{h_1 h_2} := \mathcal{P}_{h_1}^\dagger \cdot \mathcal{P}_{h_2}, \quad \widetilde{\mathcal{M}}_{h_1 h_2} := \widetilde{\mathcal{A}}_{h_1}^\dagger \cdot \widetilde{\mathcal{A}}_{h_2}, \quad (2.2.12)$$

with \cdot denoting contraction of colour indices. The off-diagonal terms ($h_1 \neq h_2$) are called spin correlations and make subdominant contributions to the sum. They are omitted in the leading-spin (LS) approximation. Particularly for high-multiplicity processes, the spin correlations can become significant deep in the IR limit, in which case a full-spin (FS) evaluation is necessary.

2.2.2 Soft limits

Next, we consider the tree-level limit when a single parton goes soft. This is only possible with a gluon, *i.e.* the limit \mathcal{E}_g , as taking a quark in the exact soft limit would result in a violation of quark number conservation in the reduced amplitude. The full $(n+1)$ -parton amplitude factorises to an n -parton reduced amplitude and an eikonal amplitude with dependence on three legs of the full amplitude,

$$\mathcal{A}_{n+1}^{(0)} \rightarrow \widetilde{\mathcal{A}}_n^{(0)} \otimes \mathcal{S}_1^{(0)}. \quad (2.2.13)$$

In the limit, the $(n+1)$ -parton phase space reduces to an n -parton phase space,

$$\{\dots, p_i, \dots, p_j, \dots, p_k, \dots\}_{n+1} \xrightarrow{E_j \rightarrow 0} \{\dots, \tilde{p}_i, \dots, \tilde{p}_k, \dots\}_n, \quad (2.2.14)$$

where j is the soft leg. Near the limit, we can map the full momenta to the reduced momenta while recoiling to enforce momentum conservation and on-shellness in the reduced phase space, for example,

$$\begin{aligned} \tilde{p}_i &= p_i + p_j - \beta p_k, & \beta &= \frac{s_{ij}}{s_{ik} + s_{jk}}. \\ \tilde{p}_k &= (1 + \beta) p_k, \end{aligned} \quad (2.2.15)$$

The kinematic part of the eikonal amplitude takes a simple form,

$$S_+(i, j, k) = \frac{\langle ik \rangle}{\langle ij \rangle \langle jk \rangle}, \quad S_-(i, j, k) = -\frac{[ik]}{[ij][jk]}, \quad (2.2.16)$$

where the subscript denotes the helicity of the soft leg j . These expressions are universal: they are independent of the parton type and helicity of partons i and k . This limit can be pictorially represented for a partial amplitude as

$$\text{Diagram} \xrightarrow{E_j \rightarrow 0} \text{Diagram} \times \text{Diagram}, \quad (2.2.17)$$

where p_i emits the soft p_j , with p_k as a reference momentum.

While kinematics factorise exactly at amplitude level, no matter how soft the gluon is, it still carries the same colour charge. Thus, colour correlations are produced. In fact, to reproduce the full colour structure of an $(n+1)$ -parton colour-summed amplitude from the reduced and eikonal amplitudes, it is necessary to sum over emissions of the soft gluon from all the external legs of the reduced amplitude,

$$\mathcal{A} \xrightarrow{E_j \rightarrow 0} \sqrt{2} \sum_{\substack{i=1 \\ i \neq j}}^{n+1} \tilde{\mathcal{C}}^{(i)} S_{h_j}(i, j, k), \quad \tilde{\mathcal{C}}^{(i)} := \tilde{R} \tilde{\mathcal{C}} \cdot T^{(i)}, \quad (2.2.18)$$

where we adopt a brief notation to show the colour-correlated amplitude as $\tilde{\mathcal{C}}^{(i)}$, the reduced amplitude decomposed into its kinematics \tilde{R} and colour $\tilde{\mathcal{C}}$, and the eikonal colour contribution as $T^{(i)}$. The \cdot represents the partial colour contraction of “attaching” the soft gluon j to leg i . The parton type of leg i determines the form of $T^{(i)}$:

$$T^{(i)} = \begin{cases} f^{a_i a_j a_k} & \text{gluon,} \\ t_{f_i f_k}^{a_j} & \text{quark,} \\ \bar{t}_{f_i \bar{f}_k}^{a_j} = -t_{f_k f_i}^{a_j} & \text{antiquark.} \end{cases} \quad (2.2.19)$$

Note that the sum in Eq. (2.2.18) runs over the indices of the reduced n -parton phase space. Since for each term in the sum i is different and k may be different, a local mapping Eq. (2.2.14) is required for each term. This is in contrast to the global mapping Eq. (2.2.4) we found for the collinear limit.

For the squared helicity amplitude,

$$|\mathcal{A}|^2 \xrightarrow{E_j \rightarrow 0} \sum_{i,k} \tilde{\mathcal{C}}^{(i,k)} S_{j^{h_j}}^{(i,k)}, \quad (2.2.20)$$

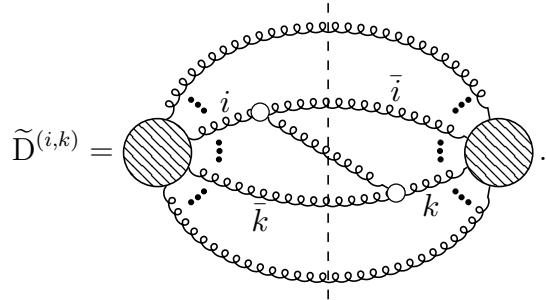
with

$$\tilde{\mathcal{C}}^{(i,k)} = \tilde{\mathcal{C}}^{(i)\dagger} \cdot \tilde{\mathcal{C}}^{(k)}, \quad S_{j^{h_j}}^{(i,k)} = S_{h_j}(i, j, k)^\dagger S_{h_j}(k, j, i) = -2 \frac{s_{ik}}{s_{ij} s_{jk}}, \quad (2.2.21)$$

where i and k run over the indices of the reduced phase space. The colour-correlated reduced ME $\tilde{C}^{(i,k)}$ can be constructed from the reduced partial amplitude vector and the colour correlations as

$$\begin{aligned} \tilde{C}^{(i,k)} &= \tilde{R}_a^\dagger \left(\tilde{D}^{(i,k)} \right)^{ab} \tilde{R}_b, \\ \left(\tilde{D}^{(i,k)} \right)^{ab} &:= \tilde{C}^a_{\dots c_i \dots c_{\bar{k}} \dots} T_{c_i c_j c_{\bar{i}}}^{(i)} T_{c_k c_j c_{\bar{k}}}^{(k)} \tilde{C}^b_{\dots c_k \dots c_{\bar{i}} \dots}, \end{aligned} \quad (2.2.22)$$

where a and b run over the colour basis. In the definition of the colour-correlation matrix $\left(\tilde{D}^{(i,k)} \right)^{ab}$ we explicitly show abstract colour indices, with remaining contractions occurring over the omitted indices. Expressed in diagram form, the colour-correlation matrix is



$$\tilde{D}^{(i,k)} = \text{Diagram} \quad (2.2.23)$$

As the degree of colour correlation grows with the number of legs on the amplitude, construction of the reduced ME $\tilde{C}^{(i,k)}$ can be computationally intensive at high multiplicities. Notice that the diagonal of the eikonal matrix vanishes by masslessness and the reduced ME is symmetric,

$$S_{j^{h_j}}^{(i,i)} = 0, \quad \tilde{C}^{(i,k)} = \tilde{C}^{(k,i)}, \quad (2.2.24)$$

so Eq. (2.2.20) can be optimised as

$$|\mathcal{A}|^2 \xrightarrow{E_j \rightarrow 0} 2 \sum_{\substack{i,k \\ k>i}} \tilde{C}^{(i,k)} S_{j^{h_j}}^{(i,k)}. \quad (2.2.25)$$

In contrast to the notation used here, the literature generally treats the colour with a notation described in Section 3.1 of Ref. [221]. This notation uses an abstract colour charge \mathbf{T}_i which comes with the emission of a gluon from the i^{th} parton. The limit expression is constructed from an n -parton reduced ME and an eikonal current,

$$\mathbf{J}^\mu(q) = \sum_{i=1}^n \mathbf{T}_i \frac{p_i^\mu}{p_i \cdot q}. \quad (2.2.26)$$

To see the connection with our notation, the eikonal amplitudes, Eq. (2.2.16), can be obtained from the kinematic part of the eikonal current by contracting with a gluon polarisation vector of the appropriate helicity, Eq. (1.5.46), and using Eq. (1.5.37).

2.2.3 Beyond NLO

The previously considered limits, double-collinear and single-soft, are single-unresolved configurations and first appear in the real corrections at NLO. For RR contributions at NNLO, we also have double-unresolved configurations, which factorise \mathcal{A}_{n+2} to \mathcal{A}_n [221, 222]. This involves multiple soft or collinear pairs of particles, or a mixture of the two, which includes: correlated limits, which introduce new singular functions; and uncorrelated limits, where the reduced amplitude is iteratively factorised by taking multiple single-unresolved limits.

One correlated double-unresolved limit is the triple-collinear limit, which up to antiparticles and permutations includes $g \rightarrow ggg$, $q \rightarrow qgg$ (with $g \rightarrow \bar{q}qg$ by crossing symmetry), and $q \rightarrow q\bar{Q}Q$, where Q may be a quark of different flavour to q . The amplitude factorises as

$$\mathcal{A}_{n+2}^{(0)} \rightarrow \tilde{\mathcal{A}}_n^{(0)} \otimes \mathcal{P}_2^{(0)}. \quad (2.2.27)$$

For instance, $g \rightarrow ggg$ can be drawn as

$$\text{Diagrammatic equation (2.2.28)} \quad (2.2.28)$$

This is treated similarly to the double-collinear limit. In analogy to the parametrisation of Eqs. (2.2.3a) to (2.2.3d), we can define a multi-collinear parametrisation of the kinematics [223].

When considering the RV contribution at NNLO, we need to calculate single-unresolved limits of loop-level amplitudes. With loops, we must take care to include all contributions at the fixed order. For instance, the one-loop single-unresolved limit factorises as

$$\mathcal{A}_{n+1}^{(1)} \rightarrow \tilde{\mathcal{A}}_n^{(1)} \otimes \mathcal{S}_1^{(0)} + \tilde{\mathcal{A}}_n^{(0)} \otimes \mathcal{S}_1^{(1)}. \quad (2.2.29)$$

Explicitly for the one-loop double-collinear limit [224], we have

$$\begin{aligned}
 & \text{Diagram} \xrightarrow{i||j} \sum_h \text{Diagram}_h \otimes \text{Diagram}_{-h} \\
 & + \sum_h \text{Diagram}'_h \otimes \text{Diagram}''_{-h} .
 \end{aligned}
 \tag{2.2.30}$$

2.2.4 Utility

These limits offer a convenient way to validate new higher-order expressions. In IR regions of phase space, a numerical evaluation of the full amplitude is approximated by the product of the appropriate singular function and lower-multiplicity reduced amplitude. A momentum mapping scheme may be necessary to relate the full and reduced phase space, as discussed for the double-collinear limit (Section 2.2.1). Generally, the reduced amplitude is available, or more easily calculated than the full amplitude, and the singular functions are well known in the literature. Thus, by numerically evaluating both the full and factorised expressions at points in IR limits, we can check that a new higher-order amplitude correctly displays the expected IR behaviour.

Factorised expressions can also be used to provide an implementation of amplitudes that, while valid only near their IR limit, is more numerically stable than the full amplitude [225]. The limit functions will also provide the building blocks of counterterms to regulate IR divergences at NNLO, as discussed in Section 2.3.

2.3 Infrared subtraction

Recall the perturbative expansion in α_s of the cross section discussed in Section 1.8.1. The NLO correction to the cross section for a process with n final state partons is given by

$$d\sigma_{\text{NLO}} = \int_n d\sigma_{\text{NLO}}^{\text{V}} + \int_{n+1} d\sigma_{\text{NLO}}^{\text{R}}, \tag{2.3.1}$$

where $d\sigma_{\text{NLO}}^{\text{V}}$ and $d\sigma_{\text{NLO}}^{\text{R}}$ indicate respectively the renormalised virtual and the real correction, with the real integration over a higher-dimensional phase space due to the additional emission. The virtual amplitude contains IR divergent loops

(Section 1.3.2), while the real cross section diverges in single-unresolved corners of phase space.

While possible in principle, analytical integration in $4 - 2\epsilon$ dimensions of real-type corrections to cross sections for modern phenomenology is computationally intractable. In addition, non-perturbative effects in simulations are generally computed numerically, so analytic integration is not possible. Instead, we use the numerical technique of MC integration (Section 1.8.2), which requires an integer number of dimensions. The method of IR subtraction is used to regulate IR divergences while allowing efficient integration.

The aim of IR subtraction is to separately cancel the divergences in the virtual and real contributions. We introduce a real subtraction term $d\sigma_{\text{NLO}}^{\text{S}}$ which locally isolates the singular behaviour of the real correction. This term can be constructed by exploiting the previously discussed factorisation properties of QCD in IR limits. It is important for computational efficiency that the subtraction term is as simple as possible while still reproducing the correct IR behaviour. The divergences of the virtual correction are subtracted at the level of the amplitude to define the finite remainders (FRs), as in Eq. (4.2.7); we represent this schematically as a virtual subtraction term $d\sigma_{\text{NLO}}^{\text{T}}$. The NLO correction can then be rewritten as

$$d\sigma_{\text{NLO}} = \int_n \left(d\sigma_{\text{NLO}}^{\text{V}} - d\sigma_{\text{NLO}}^{\text{T}} \right) + \int_{n+1} \left(d\sigma_{\text{NLO}}^{\text{R}} - d\sigma_{\text{NLO}}^{\text{S}} \right), \quad (2.3.2)$$

with

$$d\sigma_{\text{NLO}}^{\text{T}} + \int_1 d\sigma_{\text{NLO}}^{\text{S}} = \mathcal{O}(\epsilon), \quad (2.3.3)$$

such that both integrals in Eq. (2.3.2) are IR finite in $\epsilon \rightarrow 0$ and can therefore be numerically evaluated through MC integration. Various definitions of the subtraction terms give rise to different subtraction schemes. Depending on the scheme, the real subtraction term can be required to be analytically integrable over the phase space of the unresolved parton, such that the matching of the singularities of the virtual corrections in Eq. (2.3.3) occurs in a fully analytical fashion.

At NLO, general algorithms are firmly established for IR subtraction, including Frixione-Kunszt-Signer (FKS) subtraction [226, 227] and CS subtraction [228, 229]. The implementation of the subtraction terms at NNLO is not yet automated to this degree [230–232]. However, many regularisation schemes are proposed, including:

- antenna subtraction [233–236] (we use this scheme to regulate the cross section in Section 5.2),
- sector-improved residue subtraction [237–239],

- nested soft-collinear subtraction [240],
- CoLoRFulNNLO subtraction [241] (closest to CS subtraction),
- projection-to-Born subtraction [242],
- local analytic sector subtraction (Torino scheme) [243, 244],
- q_T -slicing [245],
- n-jettiness slicing [246],
- geometric IR slicing [247].

2.4 Implementation

In version 3.0.0 of `NJet3` [1], we present an analytic library of the various soft and collinear limit helicity functions of QCD amplitudes. We use amplitude-level factorisation [248] for efficient construction of factorised MEs including spin correlations.

It is possible that 64-bit floating-point numbers (`f64s`) do not provide sufficient precision when probing deep in the limit. For this case, all `NJet3` classes are templated to support higher fixed precisions of 128-bit floating-point number (`f128`) and 256-bit floating-point number (`f256`) provided by the `QD` library [249].

We make available the splitting amplitudes \mathcal{P}_h , spin matrices $\mathcal{P}_{h_1 h_2}$, and colour-correlation matrices $\tilde{D}^{(i,j)}$ for the following partonic limits:

- the single-soft tree-level limit, \mathcal{E}_g ,
- all independent double-collinear limits up to one-loop, $g \rightarrow gg$ and $q \rightarrow qq$, from [250],
- the triple-collinear tree-level limits, $g \rightarrow ggg$ and $g \rightarrow \bar{q}qg$, from [223].

As amplitudes with photons and gluons can be constructed from pure-gluon amplitudes by summing over permutations [251, 252], the loop-induced mixed QCD+QED limit functions, $g \rightarrow \gamma g$ and $g \rightarrow \gamma\gamma$, can also be generated from this set.

This includes all NLO limits, and all limits for RV at NNLO, but omits the $q \rightarrow q\bar{Q}Q$ [253] and double-soft limits that appear in RR. While therefore currently incomplete for NNLO, this library can be used for validation of new amplitudes and extended to provide a full library of the limits necessary to build counterterms at NNLO within a consistent framework.

2.4.1 Validation

To demonstrate the stability of the limit functions, we show the behaviour of the factorised expression compared to the full ME over a phase space path driving into the collinear limit for several cases. The full and reduced amplitudes are provided by the existing processes in the `NJet3` library. We use colour- and helicity-summed MEs for these tests.

We begin with the 3 || 4 double-collinear limit Eq. (2.2.2), considering tree-level five- and six-gluon scattering.

We parametrise phase space using the prescription of Ref. [223] and use this to generate a 100-point slice that approaches a collinear limit. The prescription uplifts an n -point phase space to an on-shell mass-conserving $(n + 1)$ -point phase space that is parametrised in a collinear limit. For the double-collinear limit, this is exactly Eq. (2.2.4) in reverse, with p_i^μ and p_j^μ given by Eqs. (2.2.3a) and (2.2.3b) respectively along with

$$q^\mu = \tilde{p}_k^\mu, \quad p_\perp^\mu = \delta \operatorname{Im}([\tilde{p} \sigma^\mu \tilde{p}_k] - [\tilde{p}_k \sigma^\mu \tilde{p}]), \quad (2.4.1)$$

and

$$p_k^\mu = \left(1 + \frac{p_\perp^2}{z(1-z)2\tilde{p} \cdot \tilde{p}_k}\right) \tilde{p}_k^\mu, \quad (2.4.2)$$

such that δ controls the degree of collinearity of the point. The exact limit is given by $\delta = 0$. We set $z = 0.5$ and vary δ over the slice as shown in Fig. 2.1. The four-point seed phase space is generated randomly using the algorithm from Ref. [254], which is provided in the `NJet3` library.

For each point in the slice, we construct the reduced phase space using the CS mapping defined by Eqs. (2.2.3c), (2.2.4) and (2.2.5). We evaluate both sides of Eq. (2.2.11) in `f64`, `f128`, and `f256` and plot their relative differences,

$$\mathcal{D}_\parallel = \frac{|\mathcal{A}|^2 - \sum_{h_1, h_2} \widetilde{\mathcal{M}}_{h_1 h_2} \mathcal{P}_{-h_1 - h_2}}{|\mathcal{A}|^2}, \quad (2.4.3)$$

in Fig. 2.1. The values of the factorised and full MEs smoothly become more similar as we probe deeper into the collinear limit. The five-gluon amplitude `f64` evaluation becomes numerically unstable below $s_{34}/s_{12} \approx 10^{-10}$, where we see \mathcal{D}_\parallel deviate from the `f128` evaluation. While the five-gluon amplitude in `NJet3` is implemented as a hard-coded analytic expression, the six-gluon amplitude is implemented numerically and becomes unstable with FS `f64` evaluation at the higher value of $s_{34}/s_{12} \approx 10^{-8}$. The FS six-gluon `f128` evaluation begins to destabilise as it approaches $s_{34}/s_{12} \approx 10^{-16}$.

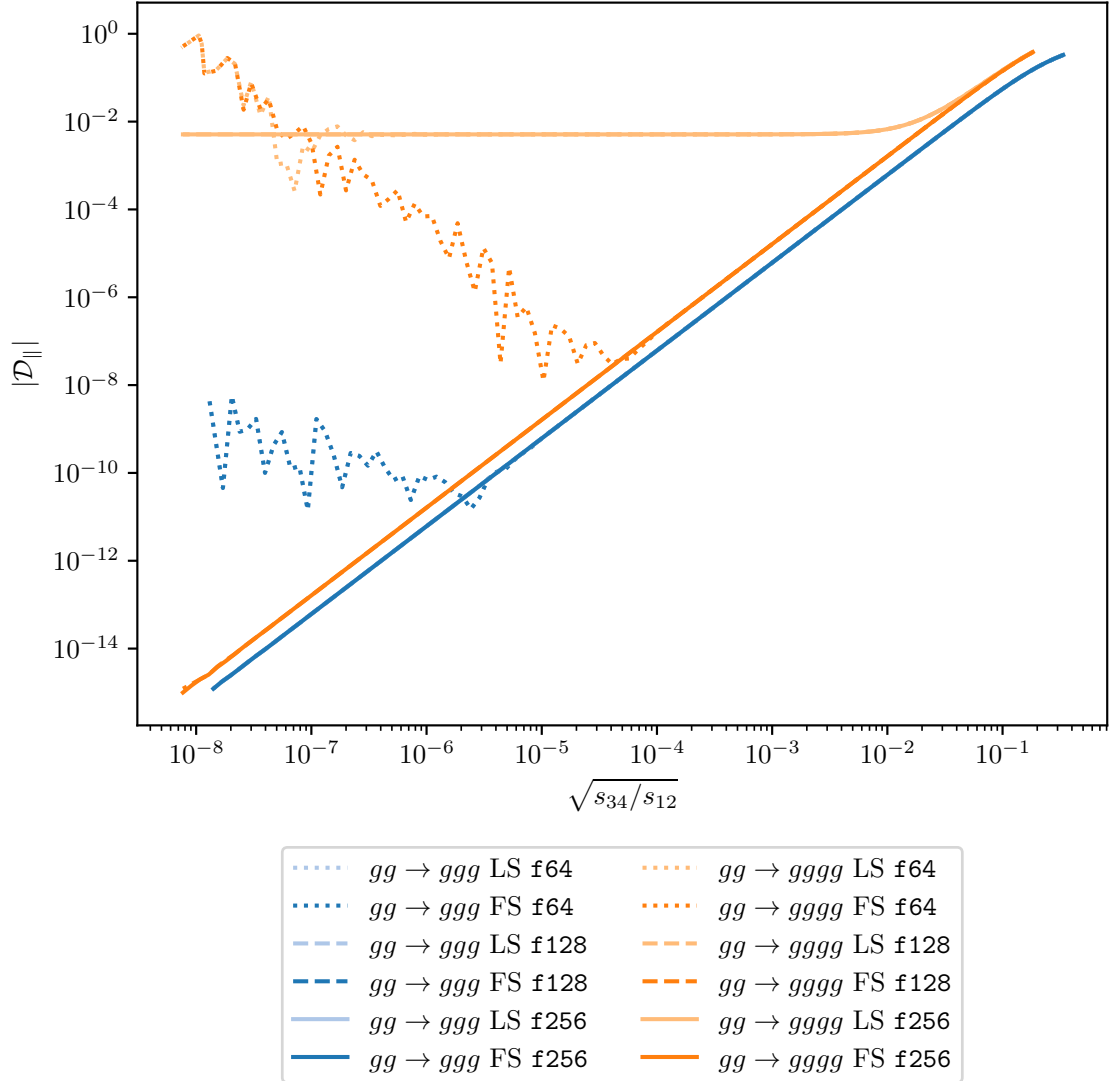


Figure 2.1: The scaling behaviour of the relative difference Eq. (2.4.3) between the factorised expression and the full ME in a collinear limit at tree level. We consider two multiplicities of gluon scattering at LS and FS with three numerical precisions.




Figure 2.2: The scaling behaviour of the relative difference between the factorised expression and the full ME in a collinear limit for the interference between the tree-level and one-loop amplitudes of $gg \rightarrow ggg$.

The five-gluon ME does not carry any spin correlations, so the LS evaluations match the FS ones. At six-gluons, however, the spin correlations are nonzero. We see the LS approximation level off at $s_{34}/s_{12} \approx 10^{-3}$, while the FS evaluation continues to approach the value of the full amplitude deeper into the limit. The LS f64 numerical instability does not become visible until the fluctuations grow to the scale of the ME, which is where the FS f64 line approaches it at $s_{34}/s_{12} \approx 10^{-13}$.

We also consider the $3 \parallel 4$ double-collinear limit Eq. (2.2.29) of five-gluon scattering at one-loop level; the full ME is the interference of the tree-level and one-loop five-gluon amplitudes. We use the same phase space parametrisation as for the tree-level case. `OneL0op` [255] is used to provide the one-loop scalar integrals. Evaluations

are performed in **f64**, **f128**, and **f256**. In Fig. 2.2, we plot the ME coefficients m_a defined by

$$\mathcal{A}^{(0)*} \mathcal{A}^{(1)} = \sum_{a=-2}^0 m_a \epsilon^a + \mathcal{O}(\epsilon). \quad (2.4.4)$$

The one-loop integrals become unstable at $s_{34}/s_{12} \approx 10^{-6}$ in **f64**, while the ϵ pole coefficients m_{-2} and m_{-1} become unstable near $s_{34}/s_{12} \approx 10^{-13}$ in **f128**. The \mathcal{D}_{\parallel} for the **f256** evaluations smoothly decreases for all coefficients, although the finite part m_0 has a much more gradual slope than the poles after $s_{34}/s_{12} \approx 10^{-3}$. Only m_0 has nonzero spin correlations, which approach from below the magnitude of the diagonal spin contributions towards the left side of the figure and are therefore not easily visible on the logarithmic scale.

2.4.2 Example code

Collinear limit

Consider the $3 \parallel 4$ double-collinear limit of a five-gluon tree-level helicity amplitude,

$$\mathcal{A}_5^{(0)}(1_g^+, 2_g^+, 3_g^+, 4_g^-, 5_g^-) \xrightarrow{3 \parallel 4} \tilde{\mathcal{A}}_4^{(0)}(1_g^+, 2_g^+, \rho_g^h, 5_g^-) \otimes \mathcal{P}_{g \rightarrow gg}^{(0)}(3^+, \rho_g^{-h}, 4^-). \quad (2.4.5)$$

To demonstrate the use of the spin matrices $\mathcal{P}_{h_1 h_2}$ in **NJet3**, we show how to evaluate Eq. (2.2.11) for this case in the following **C++** code.

```

1  #include <algorithm>
2  #include <array>
3  #include <cmath>
4  #include <complex>
5  #include <iomanip>
6  #include <iostream>
7
8  #include "analytic/0q4g-analytic.h"
9  #include "analytic/0q5g-analytic.h"
10 #include "ir/split_g2gg-analytic.h"
11 #include "ngluon2/Mom.h"
12
13 int main()
14 {
15     // 2 || 3 collinear massless 5-particle phase space (0 indexing)
16     const std::array<MOM<double>, 5> full_mom { {
17         { -5.0000000000000000e-01, 0.0000000000000000e+00,
18         ↪ 0.0000000000000000e+00, 5.0000000000000000e-01 },

```

```

18     { 4.9999949999999999e-01, 4.3301226887951738e-01,
    ↪ 0.0000000000000000e+00, 2.49999750000000005e-01 },
19     { 2.9999999999999999e-01, -2.5960910680884997e-01,
    ↪ 2.8795483728213630e-04, -1.5034303689869646e-01 },
20     { 2.00000500000000008e-01, -1.7340316207066742e-01,
    ↪ -2.8795483728213630e-04, -9.9656713101303585e-02 },
21     { -5.0000000000000000e-01, 0.0000000000000000e+00,
    ↪ 0.0000000000000000e+00, -5.0000000000000000e-01 },
22 } };
23
24 // Catani-Seymour momentum mapping for 2 || 3 (0 indexing)
25 const MOM<double> p { full_mom[2] + full_mom[3] };
26 const double x { dot(full_mom[2], full_mom[3]) / dot(p, full_mom[4])
    ↪ };
27
28 // reduced phase space
29 std::array<MOM<double>, 4> reduced_mom {};
30 std::copy_n(full_mom.cbegin(), 2, reduced_mom.begin());
31 reduced_mom[2] = p - x * full_mom[4];
32 reduced_mom[3] = (1. + x) * full_mom[4];
33
34 // momenta for splitting amplitude
35 const std::array<MOM<double>, 3> splitting_mom { { -reduced_mom[2],
    ↪ full_mom[2], full_mom[3] } };
36
37 // initialise amplitudes
38 Amp0q5g_a<double> full_amp;
39 Amp0q4g_a<double> reduced_amp;
40 Splitg2gg_a<double> splitting_amp;
41
42 // set momenta
43 full_amp.setMomenta(full_mom.data());
44 reduced_amp.setMomenta(reduced_mom.data());
45 // the second argument is a reference momentum
46 splitting_amp.setMomenta(splitting_mom.data(), reduced_mom[3]);
47
48 // helicities for amplitudes
49 const std::array<int, 5> full_hels { +1, +1, +1, -1, -1 };
50 // the helicity values of the correlated legs do not matter
51 const std::array<int, 4> reduced_hels { +1, +1, 0, -1 };

```

```

52     const std::array<int, 3> splitting_hels { 0, +1, -1 };
53
54     // initialise (flattened) 2x2 spin matrices
55     // order: ++ +- -+ --
56     std::array<std::complex<double>, 4> reduced_spn_mat {},
57     ↪ splitting_spn_mat {};
58
59     // compute spin matrices
60     // the first argument is the index of the correlated leg
61     reduced_amp.born_spmatrix(2, reduced_hels.data(),
62     ↪ reduced_spn_mat.data());
63     splitting_amp.born_spmatrix(0, splitting_hels.data(),
64     ↪ splitting_spn_mat.data());
65
66     // while the final result will be real, the intermediate steps may be
67     ↪ complex
68     std::complex<double> limit {};
69
70     // compute factorised matrix element including spin correlations
71     for (int i { 0 }; i < 4; ++i) {
72         limit += reduced_spn_mat[i] * splitting_spn_mat[3 - i];
73     }
74
75     // compute full matrix element
76     const double amp { full_amp.born(full_hels.data()) };
77
78     // print results (with unity indexing)
79     std::cout
80     << std::setprecision(1) << std::scientific
81     << '\n'
82     << std::setw(25) << "|s_{34}/s_{12}| = "
83     << std::setw(7) << std::abs(dot(full_mom[2], full_mom[3]) /
84     ↪ dot(full_mom[0], full_mom[1]))
85     << '\n'
86     << std::setw(25) << "|(amp^2-lim^2)/amp^2| = "
87     << std::setw(7) << std::abs((amp - limit.real()) / amp)
88     << '\n'
89     << '\n';
90 }

```

This runs to give the result:

```
1      |s_{34}/s_{12}| = 1.3e-06
2      |(amp^2-lim^2)/amp^2| = 7.9e-04
```

The small ratio $s_{34}/s_{12} = 1.3 \times 10^{-6}$ shows that we are near a $3 \parallel 4$ collinear limit. The relative difference Eq. (2.4.3) between the factorised and full MEs $\mathcal{D}_{\parallel} = 7.9 \times 10^{-4}$ is small, showing a similar order of magnitude to the helicity-summed value in Fig. 2.1. This code can be used for the helicity-summed ME by simply removing the helicity argument in all function calls.

Soft limit

Consider the \mathcal{E}_4 single-soft limit of a five-gluon tree-level helicity amplitude,

$$\mathcal{A}_5^{(0)}(1_g^+, 2_g^+, 3_g^-, 4_g^-, 5_g^-) \xrightarrow{E_4 \rightarrow 0} \tilde{\mathcal{A}}_4^{(0)}(1_g^+, 2_g^+, 3_g^-, 5_g^-) \otimes \mathcal{S}_1^{(0)}(4_g^-). \quad (2.4.6)$$

To demonstrate the use of the colour-correlation matrices $\tilde{\mathbf{D}}^{(i,j)}$ in NJet3, we show how to evaluate Eq. (2.2.25) for this case in the following C++ code.

```
1  #include <algorithm>
2  #include <array>
3  #include <cmath>
4  #include <complex>
5  #include <iomanip>
6  #include <iostream>
7
8  #include "analytic/0q4g-analytic.h"
9  #include "analytic/0q5g-analytic.h"
10 #include "ir/soft_gg2g-analytic.h"
11 #include "ngluon2/Mom.h"
12
13 int main()
14 {
15     // massless 5-point phase space with full_mom[3] soft
16     const std::array<MOM<double>, 5> full_mom { {
17         { -5.0000000000000000e-01, 0.0000000000000000e+00,
18         ↪ 0.0000000000000000e+00, -5.0000000000000000e-01 },
19         { -5.0000000000000000e-01, 0.0000000000000000e+00,
20         ↪ 0.0000000000000000e+00, 5.0000000000000000e-01 },
21         { 4.9999949999999999e-01, 4.3301226887951738e-01,
22         ↪ 0.0000000000000000e+00, 2.4999975000000000e-01 },
```

```

20     { 4.9999999999772449e-06, -2.2505887718767004e-06,
    ↪ 1.7633487034433740e-06, -4.1018839000804920e-06 },
21     { 4.9999550000000009e-01, -4.3301001829074548e-01,
    ↪ -1.7633487034433740e-06, -2.4999564811609998e-01 },
22 } };
23
24 // initialise amplitudes
25 Amp0q5g_a<double> full_amp;
26 Amp0q4g_a<double> reduced_amp;
27 Softgg2g_a<double> soft_amp;
28
29 // set momenta
30 full_amp.setMomenta(full_mom.data());
31
32 // helicities for amplitudes
33 const std::array<int, 5> full_hels { +1, +1, -1, -1, -1 };
34 const std::array<int, 4> reduced_hels { +1, +1, -1, -1 };
35 // the eikonal amplitude is independent of the helicities of the
    ↪ correlated legs
36 const std::array<int, 3> soft_hels { 0, -1, 0 };
37
38 // perform reduced colour sum
39 // indices i,3,k are in the full 5-point phase space
40 // indices ii,kk are in the reduced 4-point phase space
41 double lim {};
42 int ref_index {};
43 for (int i { 0 }; i < 5; ++i) {
44     if (i != 3) {
45         const int ii { i < 3 ? i : (i - 1) % 4 };
46
47         for (int k { i + 1 }; k < 5; ++k) {
48             if (k != 3) {
49                 const int kk { k < 3 ? k : (k - 1) % 4 };
50
51                 // reduced phase space is just the full phase space
    ↪ without the soft leg
52                 std::array<MOM<double>, 4> reduced_mom {};
53                 std::copy_n(full_mom.cbegin(), 3,
    ↪ reduced_mom.begin());
54                 reduced_mom[3] = full_mom[4];

```

```

55
56         // ensure physical reduced phase space
57         const double b { dot(full_mom[i], full_mom[3]) /
↪ (dot(full_mom[i], full_mom[k]) + dot(full_mom[3], full_mom[k]))
↪ };
58         reduced_mom[ii] = full_mom[i] + full_mom[3] - b *
↪ full_mom[k];
59         reduced_mom[kk] = (1. + b) * full_mom[k];
60
61         reduced_amp.setMomenta(reduced_mom.data());
62
63         // ensure reference momentum is not the same as
↪ another momentum in use
64         // to avoid division by zero from brackets in the
↪ denominator of the form < i i >
65         while ((ref_index == i) || (ref_index == 3) ||
↪ (ref_index == k)) {
66             ref_index = (ref_index + 1) % 5;
67         }
68
69         // set momenta for eikonal amplitude
70         // the second argument is the reference momentum
71         soft_amp.setMomenta({ reduced_mom[ii], full_mom[3],
↪ reduced_mom[kk] }, full_mom[ref_index]);
72
73         // compute eikonal and colour-correlated reduced
↪ matrix elements
74         const double soft_val {
↪ soft_amp.born(soft_hels.data()) },
75         cc_val {
↪ reduced_amp.born_ccij(reduced_hels.data(), ii, kk) };
76
77         lim += cc_val * soft_val;
78     }
79 }
80 }
81 }
82
83 // we summed over upper triangle of colour correlation matrix
84 // diagonals receive no contribution because eikonal kinematic factor
↪ is zero due to antisymmetry of brackets

```

```

85 // matrix is symmetric, so add lower triangle by doubling result
86 lim *= 2.;
87
88 // compute full matrix element
89 const double amp_val { full_amp.born(full_hels.data()) };
90
91 // print results (with unity indexing)
92 std::cout
93     << std::setprecision(1) << std::scientific
94     << '\n'
95     << std::setw(25) << "|{E_4}^2/s_{12}| = "
96     << std::setw(7) << std::abs(pow(full_mom[3].x0, 2) /
→ dot(full_mom[0], full_mom[1])) << '\n'
97     << std::setw(25) << "|(amp^2-lim^2)/amp^2| = "
98     << std::setw(7) << std::abs((amp_val - lim) / amp_val) << '\n'
99     << '\n';
100 }

```

This evaluates to:

```

1 |{E_4}^2/s_{12}| = 5.0e-11
2 |(amp^2-lim^2)/amp^2| = 3.1e-05

```

The small ratio $E_4^2/s_{12} = 5.0 \times 10^{-11}$ shows that we are in an \mathcal{E}_4 soft limit. The relative difference,

$$\mathcal{D}_s = \frac{|\mathcal{A}|^2 - \sum_{i,k} \tilde{C}^{(i,k)} S_{j^{h_j}}^{(i,k)}}{|\mathcal{A}|^2}, \quad (2.4.7)$$

between the factorised and full MEs $\mathcal{D}_s = 3.1 \times 10^{-5}$ is small, supporting that the factorisation Eq. (2.2.25) holds.

Chapter 3

Matrix element neural networks

Precision phenomenological studies of high-multiplicity scattering processes present a substantial theoretical challenge and are vitally important ingredients in measurements at collider experiments. Machine learning (ML) technology has the potential to dramatically optimise simulations for complicated final states. We investigate the use of NNs to approximate MEs, studying the case of loop-induced diphoton-plus-jets production through gluon fusion. We train NN models on one-loop amplitudes from the `NJet3 C++` library [1, 256, 257] (see also Chapter 4) and interface them with the *Simulation of High Energy Reactions of Particles* (`Sherpa`) MC event generator [56, 57] to provide the ME within a realistic hadron collider simulation. Computing some standard observables, such as jet transverse momentum, with the models and comparing to conventional techniques, we find excellent agreement in the distributions and a reduced total simulation time by a factor of thirty.

This chapter is organised as follows. We first motivate the use of ME NNs in the gluon-initiated diphoton-plus-jets sector in Section 3.1. Then, in Section 3.2, we discuss the gluon-initiated diphoton-plus-jets amplitudes and their implementations within `NJet3` which form the target distribution for training the NNs. We then describe the phase-space partitioning used to handle IR divergent regions in Section 3.3. Next, in Section 3.4, we present the architecture of the NNs used. In Section 3.5, we discuss the simulation pipeline and interface of the NN model to `Sherpa`. Finally, in Section 3.6, we study the performance of the model compared to `NJet3` for $gg \rightarrow \gamma\gamma gg$ and present some distributions before concluding in Section 3.7.

Our code is publicly available [258].

3.1 Background

With the increasing size of the LHC dataset driving ever more precise experimental measurements, SM predictions for high-multiplicity scattering at hadron colliders




Figure 3.1: Diagram of $gg \rightarrow \gamma\gamma gg$ ($N = 6$) at LO. The photons couple to an internal quark loop.

form a vital part of precision phenomenology studies. These calculations mainly rely on automated numerical codes [139] to calculate amplitudes up to high multiplicities, including tree-level real corrections at NLO and RR corrections at NNLO, and one-loop RV corrections at NNLO (Section 1.8.1). These codes have been a theoretical revolution, particularly at one-loop (Section 1.6). However, due to the high dimensionality of the phase space, these real-type corrections are often the computational bottleneck in higher-order calculations (for instance, see Section 5.4).

As discussed in Section 1.8.3, the gluon-fusion channel of diphoton-plus-jets production is of high phenomenological relevance. Therefore, we study the loop-induced class of processes with two photons and many gluons (Section 1.8.1). We stress that because they are loop induced, these amplitudes are finite in ϵ (Section 1.3.2).

ML technology has found a wealth of application in high-energy physics: see the reviews [259, 260] and references therein. For an introduction to ML, see Ref. [261]. We employ the ensemble NN model of Ref. [262], which studied e^+e^- annihilation to jets, to emulate the gluon-initiated diphoton-plus-jets MEs within a full MC event generator simulation (Section 1.2.1). This tests the methodology against the additional complexity of hadron collider simulations, including PDF convolution and variable partonic centre-of-mass scale, complex phase-space cuts and jet clustering (Section 2.1.1), and phase-space sampling optimisation methods of integrators.

3.2 Amplitudes

As there is no vertex coupling gluons to photons in the SM, diphoton-plus-jets production through gluon fusion is loop induced, as depicted in Fig. 3.1 and discussed in Section 1.8.1. The LO process is $\mathcal{O}(\alpha_s^{N-2})$ for multiplicity N , appearing at NNLO in the perturbative expansion of the combined quark- and gluon-initiated process, as shown in Fig. 1.6. We study the channels with only gluons and photons in the external particles $gg \rightarrow \gamma\gamma + n \times g$. These proceed through a quark loop at LO.




Figure 3.2: Typical evaluation time of the ME for a single phase-space point. Results are shown for available implementations at various multiplicities, including numerical (blue cross) and analytical (orange triangle) evaluations using `NJet3` and inference on the NN model (green circle).

Conventional event generator simulations optimise virtual corrections in NLO calculations by learning the phase-space distribution of the LO process and using this to sample the virtual contribution. This technique fails for loop-induced processes, where the expensive one-loop amplitude has no tree-level process to optimise the phase space on. Therefore, new methods are required to improve the efficiency of integrating these channels at high multiplicity.

We use the one-loop-squared MEs from the `NJet3` library as the targets for our NN emulation. These include two classes of amplitudes: an automated numerical setup for arbitrary multiplicity; and hard-coded analytical expressions for $N \in \{4, 5\}$. The former obtains the diphoton-plus-jets amplitudes by summing permutations of pure-gluon partial amplitudes [251, 252], which are themselves based on generalised unitarity (Section 1.6.3) and integrand reduction (Section 1.6.4). While completely automated, evaluation time and numerical stability are increasingly difficult to control. The hard-coded implementations offer compact analytical expressions with extremely fast and stable evaluation, although they are unavailable for higher multiplicity. The $N = 5$ result is obtained through an FF reconstruction [167] of the permutation-sum result.

The evaluation timings of these methods are compared to the NN model in Fig. 3.2. Note that this is a single NN model, which is comprised of an FKS ensemble (Section 3.3), and not the stochastic ensemble (Section 3.5). The value is the mean of 100 evaluations using random sampling over a uniform phase space [254]. We

time single-threaded central processing unit (CPU) calls as parallelisation is applied at the level of events in event generator simulations.

3.3 Phase space partitioning

Training a single NN over the entire phase space results in a poor fit, especially at higher multiplicity [262]. This is caused by regions where the amplitude becomes IR divergent, which arise from soft (\mathcal{E}_i) and collinear ($i \parallel j$) emissions (Section 2.1). These singularities are regulated at LO with cuts (Section 2.1.1), but the amplitude in local regions exhibits extreme curvature which causes problems for the global fit. Therefore, we train a separate NN on each of the IR structures of the phase space.

We first partition the phase space into a non-divergent region $\mathcal{R}_{\text{non-div}}$ and a divergent region \mathcal{R}_{div} . Phase space points which pass the following cut are included in \mathcal{R}_{div} ,

$$\min \left(\left\{ \frac{s_{ij}}{s_{12}} : i, j \in \{1, \dots, N\} \right\} \right) < y. \quad (3.3.1)$$

The threshold y must be tuned to discriminate points of a similar scale into a single region, while having sufficient points in \mathcal{R}_{div} to train on.

We then sub-divide \mathcal{R}_{div} similarly to the decomposition of the FKS subtraction scheme (Section 2.3). We define a set of FKS pairs,

$$\mathcal{P}_{\text{FKS}} = \{(i, j) : \mathcal{E}_i \vee \mathcal{E}_j \vee i \parallel j\}, \quad (3.3.2)$$

corresponding to the singular configurations, of which there are

$$\binom{N}{2} - 1 = \frac{N^2 - N - 2}{2}. \quad (3.3.3)$$

This includes redundancy as it overcounts soft singularities, which means the model must learn this, but is favoured for its simplicity. Each pair is assigned a partition function,

$$\mathcal{S}_{ij} = \frac{1}{s_{ij} \sum_{k, \ell \in \mathcal{P}_{\text{FKS}}} \frac{1}{s_{k\ell}}}, \quad \sum_{i, j \in \mathcal{P}_{\text{FKS}}} \mathcal{S}_{ij} = 1, \quad (3.3.4)$$

which smoothly isolates that divergence on multiplication with the ME. The set of all partition functions sum to unity.

We train a NN on $|\mathcal{A}(p^\mu)|^2$ for $p^\mu \in \mathcal{R}_{\text{non-div}}$, and a NN on each of the partition-function-weighted MEs,

$$\{\mathcal{S}_{ij} |\mathcal{A}(p^\mu)|^2 : i, j \in \mathcal{P}_{\text{FKS}} ; p^\mu \in \mathcal{R}_{\text{div}}\}. \quad (3.3.5)$$




Figure 3.3: Diagram of a single NN in the $gg \rightarrow \gamma\gamma gg$ model. There are 24 input nodes (green). The hidden layer nodes are shown in blue and the output node in red.

We reconstruct the complete ME in \mathcal{R}_{div} by summing the weighted MEs,

$$|\mathcal{A}|^2 = \sum_{i,j \in \mathcal{P}_{\text{FKS}}} \mathcal{S}_{ij} |\mathcal{A}|^2. \quad (3.3.6)$$

This ensemble of NNs, referred to as the model, can be used to accurately infer the ME over the complete phase space $\mathcal{R}_{\text{non-div}} \cup \mathcal{R}_{\text{div}}$.

Note that increasing the cut y , which increases the proportion of points in \mathcal{R}_{div} , incurs a performance penalty due to the higher cost of inferring over several NNs in \mathcal{R}_{div} compared to the single NN in $\mathcal{R}_{\text{non-div}}$.

3.4 Model architecture

Although using fine-tuned architectures for each configuration (processes, cuts, etc.) would provide optimal performance, this would be prohibitively expensive in terms of personnel resources. We use a general setup as the gains of specialised NN optimisation are beyond the scope of this pioneering work, performing hyperparameter optimisation on the $gg \rightarrow \gamma\gamma gg$ process (see Appendix A of Ref. [263] for details).

As depicted in Fig. 3.3, each NN uses a fully-connected architecture, a standard

choice for a regression problem, parametrised using the `Keras Python` application programming interface (API) [264] to the `TensorFlow` ML library [265]. There are $4 \times N$ input nodes: one for each component of each momentum in the phase-space point. The three hidden layers are comprised of 20, 40, and 20 nodes respectively, all with hyperbolic-tangent activation functions, which we found to outperform standard alternatives such as rectified linear unit. There is a single output node with a linear activation function, which returns the approximation of the ME. We find that this is a sufficient number of layers and nodes to learn the MEs, while remaining economical for computational performance.

We train with a mean squared error (MSE) loss function,

$$L = \frac{1}{n} \sum_{i=1}^n (f(x_i) - y_i)^2, \quad (3.4.1)$$

where the model is represented by a function f acting on n input data points x_i with targets y_i , using Adam-optimised gradient descent [266]. We expect that the model will learn the mean of the target distribution using this loss function (Appendix A of Ref. [262]). The number of training epochs is determined by Early Stopping regularisation (Section 8.1.2 of Ref. [261]), with a patience of 100 epochs to mitigate the effects of the limited size of \mathcal{R}_{div} that may appear in the validation set. We use 32-bit floating-point numbers (`f32s`) throughout.

3.5 Pipeline

Our ML pipeline used to produce the $gg \rightarrow \gamma\gamma gg$ results presented, sketched in Fig. 3.4, is:

1. Generate the training and validation datasets by running `Sherpa` with `NJet3` on a unit integration grid, *i.e.* uniformly sampling the phase space such that all phase space weights are one. Similarly, generate the testing dataset with a different random seed.
2. Train the model.
3. Use the model to estimate the MEs during event generation with `Sherpa`, using the same integration grid.

Input data consists of a list of phase-space points $p_i^\mu \in \mathbb{R}^{4N}$ and the corresponding colour- and helicity-summed one-loop-squared ME $|\mathcal{A}|^2 \in \mathbb{R}$. Phase space sampling is determined by the integrator, meaning the training is optimal only for a specific integrator. The results presented here use the `RAMBO` integrator [267], although we also study `VEGAS` [268] in Ref. [263]. The data is extracted from a run of the

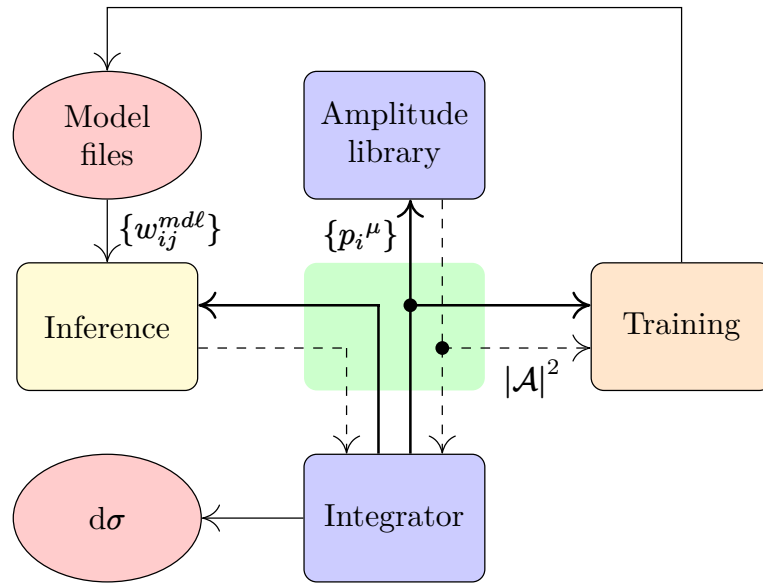


Figure 3.4: Flow chart of our ML pipeline. An MC event generator and ME library (blue) are used to generate a training dataset of phase-space points (bold line) and MEs (dashed line). Our interface code (green) extracts this data and makes it available for training (orange). The NN weights are denoted by w_{ij}^{mdl} , where m is the index of the model in the ensemble, d is the NN in the model, l denotes the layers for a link between nodes in layers l and $l+1$, and i and j identify the nodes in these layers. They are encoded in model files (red) that are read by our inference code (yellow) to provide an approximation of the ME for a given phase-space point, which can be provided by the integrator in a hadronic simulation to efficiently compute a cross section (red).

integrator, generating 100 k points which are split 4:1 into training and validation datasets. A 3 M point testing dataset is produced by a second run of the integrator with a different random number seed and used to evaluate model performance.

We perform inference on an ensemble of twenty models, each of which has different random weight initialisation and shuffled training and validation datasets. We take as the result the mean of the ensemble,

$$\bar{x} = \frac{\sum_{i=1}^{20} x_i}{20}, \quad (3.5.1)$$

where x_i is the result from each model, with the standard error providing the precision/optimalty error [262],

$$\varepsilon_{\text{std}} = \frac{\sqrt{\sum_{i=1}^{20} (x_i - \bar{x})^2}}{20}. \quad (3.5.2)$$

While training was performed using `Python`, event generators are generally written in `C++`. To use the model within a simulation, we wrote a `C++` inference code as well as a bespoke `C++` interface for `Sherpa`. The weights of the trained models are written to file and read by the inference code at runtime; the library `Eigen3` [269] is used to perform efficient linear algebra on the CPU. The interface can also be used to call `C++` amplitude libraries directly instead of the model. We use this to interface `NJet3` to `Sherpa` to generate the datasets, which is performed with `f64s`. Calls are made though a *Binoth Les Houches Accord* (BLHA) interface [270, 271], which is compatible with all BLHA-supporting amplitude libraries with minor modifications.

PDFs are provided by LHAPDF [272] using the NLO NNPDF3.1 set¹ with $\alpha_s(m_Z) = 0.118$ [273]. Cuts are adapted from those in Ref. [274]. Analysis, including all treatment of MC errors, is performed using `Rivet` [275] with an adapted reference analysis script [276].

3.6 Results

Comparing the output of the trained $gg \rightarrow \gamma\gamma gg$ model to the amplitude library value through point-by-point ratio in Fig. 3.5, we see a peaked and approximately symmetric error distribution with a shifted mean in both regions. Both region histograms have a similar mean, indicating comparable accuracy, with $\mathcal{R}_{\text{non-div}}$ performing slightly better. The distributions are fairly broad. $\mathcal{R}_{\text{non-div}}$ shows a slight tail on the right, which arises from points near the cutoff y .

Despite the per-point agreement being somewhat poor, the total cross section is

¹NNPDF31_nlo_as_0118




Figure 3.5: Histogram by region of the decimal logarithm of the ratio between the ME returned by the model and `NJet3` for each point in a 1 M subset of the testing data for $gg \rightarrow \gamma\gamma g$. The region cut is $y = 10^{-3}$ and \mathcal{R}_{div} (blue) contains 2.4% of the points.




Figure 3.6: Histogram of the root mean squared error (RMSE) of the model compared to `NJet3` for a slice of phase space in x_2 , the momentum fraction of the second incoming parton (gluon), for $gg \rightarrow \gamma\gamma g$ (yellow). Also shown are the points in the training dataset, binned in x_2 (blue).




Figure 3.7: As Fig. 3.5, but weighted by the PDFs.

found to be in agreement, with

$$\begin{aligned}\sigma_{\text{NJet3}} &= (49 \pm 5) \times 10^{-7} \text{ pb} \quad (\text{MC error}), \\ \sigma_{\text{NN}} &= (45 \pm 6) \times 10^{-7} \text{ pb} \quad (\text{precision/optimalty error}).\end{aligned}$$

Figure 3.6 shows the RMSE $\varepsilon_{\text{RMSE}}$ of the model compared to NJet3 against the frequency of points appearing on a univariate slice of phase space in the training dataset. For a bin with n points, the RMSE is calculated as

$$\varepsilon_{\text{RMSE}} = \sqrt{\frac{1}{n} \sum_{i=1}^n (r_{\text{NN}} - r_{\text{NJet3}})^2}, \quad (3.6.1)$$

for target result r_{NJet3} and model result r_{NN} . Sampling frequency during unit-grid integration is determined by the gluon PDF. The figure shows that the regions that are sampled the most due to the shape of the gluon PDF are those that have the lowest error, which is why the agreement in the total cross section is much better than for point-by-point. To validate this observation, in Fig. 3.7 we weight the per-point ratios of the model result compared to the target by the PDFs, using bin contributions for each point of

$$f_g(x_1, \mu_F) f_g(x_2, \mu_F), \quad (3.6.2)$$

where $f_g(x_i, \mu_F)$ is the gluon PDF for the momentum fraction x_i of gluon i and factorisation scale μ_F . We plot the normalised weighted histograms. The result is indeed more narrowly peaked and closer to being unit-centred, further suggesting that poorly performing points fall in PDF-suppressed regions. This indicates that

the accuracy of distributions inferred with the model is dependent on the choice of process, cuts, and observable.

Fig. 3.8 shows excellent agreement between the distributions obtained from the model and `NJet3` for two differential slices of phase space. There are some fluctuations in the tails although they appear statistical rather than systematic and the model predictions mostly remain within the `NJet3` MC error bands. Normalised NN uncertainties are negligible compared to the MC error.

In Ref. [263], we also demonstrate how agreement can be improved in \mathcal{R}_{div} by reweighting event weights by the ratio of the emulated and true MEs at known points from the training data, as well as showing good performance for $gg \rightarrow \gamma\gamma g$ when relaxing cuts at inference compared to training.

Subsequent to this work, the authors of Ref. [277] achieve improved per-point agreement at tree-level by exploiting the factorisation properties of MEs (Section 2.2).

3.7 Summary

We extend previous work which pioneered the emulation of MEs with NNs, studying these techniques for the first time within a full hadron collider simulation. We focus on loop-induced diphoton-plus-jets production via gluon fusion. The difficulties introduced by IR behaviour are tamed by partitioning the phase space as prescribed by FKS subtraction. We provide a general interface for trained models to `Sherpa`.

We find that ME NN models provide an efficient general framework for optimising high-multiplicity observables at hadron colliders. Agreement in differential distributions is excellent. As the cost of inference on the model is negligible compared to the amplitude library call in training, the speedup factor in total simulation time (including training) compared to conventional methods is given by the ratio of the number of points used for inference and training,

$$\frac{N_{\text{infer}}}{N_{\text{train}}}. \quad (3.7.1)$$

For this study, this gave a factor of thirty, although for studies with higher statistics or coverage of multiple cut configurations, the factor would be much greater.




Figure 3.8: Differential normalised cross sections for $gg \rightarrow \gamma\gamma gg$, comparing NJet3 (red; MC error) to the model (blue; precision/optimalty error), in R-separation, Eq. (2.1.4), between the hardest jet and photon (upper) and the transverse momentum of the hardest jet (lower). Refer to Ref. [263] for details of cuts and for further distributions.

Chapter 4

Virtual QCD corrections to the $gg \rightarrow \gamma\gamma g$ amplitude

We present an analytic computation of the amplitudes for the gluon-initiated contribution to diphoton-plus-jet production at hadron colliders at up to two loops in QCD. We reconstruct the analytic form of the FRs from numerical evaluations over FFs including all colour contributions. Compact expressions are found using the pentagon function basis [278]. We provide a fast and stable implementation for the colour- and helicity-summed FRs, including the one-loop squared and the interference between one- and two-loop FRs, in C++ as part of the version 3.0.0 release of the NJet3 library [1].

This chapter is organised as follows. We first cover recent progress in two-to-three two-loop amplitude computations in Section 4.1. Then, in Section 4.2, we introduce the notation and describe the colour decomposition of the amplitudes. We describe the methodology used to perform the IBP reduction in Section 4.3. In Section 4.4, we include some details of the momentum twistor formalism used to provide a rational parametrisation of the kinematics. Next, in Section 4.5, we discuss the reconstruction of the FRs over FFs. In particular, we describe a method for performing a univariate partial fractioning of the rational coefficients of the special functions on the fly in Section 4.5.3. This approach can be used inside the FF workflow, reducing significantly the number of sample points required to complete the analytic reconstruction and yielding compact analytic expressions. In Section 4.6, we show the simple analytic forms we obtained for the all-plus helicity amplitude, *i.e.* all external particles with positive helicity, which highlight its conformal properties. Finally, in Section 4.7, we present the implementation in the NJet3 library and the performance of the code using a realistic set of phase-space points before concluding with a few remarks on future applications of the results and methods in Section 4.8.

4.1 Background

As discussed in Section 1.8.3, diphoton-plus-jet production is an important experimental signature. The Born-level amplitude for the gluon-initiated subprocess contains a closed quark loop coupling to both photons (Section 1.8.1). Consequently, this subprocess starts to contribute to the cross section only from NNLO onwards (Fig. 1.6). Here, we derive the two-loop virtual amplitudes for $gg \rightarrow \gamma\gamma g$ that contribute to the NLO corrections to the closed quark loop subprocess. Curiously, the gluon channel has the opposite structure to the conventional expansion in the number of colour charges N_c . The LC contributions to the quark-initiated process, which are the dominant contributions, contain only planar diagrams [210]. However, in the gluon-initiated case, the LC limit contains both planar and non-planar graphs at two loops. Graphs with the highest complexity are thus unavoidable.

The last few years have seen rapid progress in our ability to compute two-loop two-to-three scattering processes in QCD that had been intractable for a long time. The analytic computation of the scattering amplitudes in a form suitable for phenomenological applications requires overcoming a number of major technical bottlenecks. A basis of special functions must be identified that can be evaluated efficiently over the full phase space. For massless five-point scattering, such a basis has been identified [279–283] and became recently available as a fast and stable implementation in C++ valid in the physical scattering region [278]. Secondly, the amplitude must be reduced from tensor Feynman integrals onto a basis of MIs that can subsequently be expanded in terms of special functions. Currently, the only viable approach to this task is through the solution of enormous systems of IBP identities [160, 162, 163], for which many public implementations now exist [168, 170, 284–287]. There has been success in simplifying this problem using syzygy relations [187, 288–291], module intersection [292, 293], intersection theory [294–297], η expansion [298–302], direct solution of IBPs through recursive relations [303], multivariate partial fractioning [293], and by-passing complicated algebraic steps through FF arithmetic (Section 1.7). The latter method can be applied more broadly [165, 167], in particular to a complete reduction of the amplitudes into a representation using special functions. New efficient reconstruction techniques [129, 304–307] allow compact expressions of the rational coefficients to be found. In this chapter, we approach the problem through a direct analytic reconstruction of the amplitudes at the level of the pentagon functions, performing all intermediate steps numerically over FFs. This technique has been applied successfully to LC (planar) five-parton amplitudes first numerically [308–311] and then analytically [312–316] (see Chapter 6). LC triphoton production has also been completed and cross checked by two independent groups both at the level of the amplitudes [317, 318] and of differential cross sections [319, 320]. Very recently,

NNLO QCD predictions for a number of trijet observables and differential three-to-two jet ratios have been computed at LC as well [236, 321]. The process $gg \rightarrow g\gamma\gamma$ contains the most complicated non-planar topologies with up to rank five tensor numerators even at LC.

We obtain compact analytic expressions for the complete set of helicity amplitudes for which the UV and IR poles have been subtracted, and implement them into an efficient and stable C++ code as part of the `NJet3` library [1]. These expressions take the form of rational coefficients multiplied by pentagon functions. The code provides colour- and helicity-summed expressions for the interference between the one- and two-loop amplitudes, which can be used directly in phenomenological applications as in Chapter 5.

4.2 Kinematics and amplitude conventions

We consider the production of a pair of photons in association with a gluon from gluon fusion,

$$g(-p_1) + g(-p_2) \rightarrow g(p_3) + \gamma(p_4) + \gamma(p_5), \quad (4.2.1)$$

up to two-loop order in QCD. All particles are massless, Eq. (1.5.4), and we take all momenta as outgoing, Eq. (1.5.5). Without loss of generality, we assume that the external momenta p_i live in a four-dimensional Minkowski spacetime, whereas the Feynman loop integrations are done in $d = 4 - 2\epsilon$ dimensions to regulate the divergences (Section 1.3.2). As discussed in Section 1.5.1, the kinematics are described by five LI scalar invariants, which can be chosen as the set of momentum invariants $\{s_{12}, s_{23}, s_{34}, s_{45}, s_{15}\}$, and a pseudoscalar invariant tr_5 defined in Eq. (1.5.14).

We work in the s_{12} physical scattering region, which is delimited by the requirements that all s -channel invariants are positive and all t -channel invariants are negative,

$$s_{12}, s_{34}, s_{35}, s_{45} > 0, \quad (4.2.2)$$

$$s_{13}, s_{14}, s_{15}, s_{23}, s_{24}, s_{25} < 0, \quad (4.2.3)$$

together with the negativity of the Gram determinant defined in Eq. (1.5.19), $\Delta < 0$, which follows from the real-valuedness of the momenta [281].

The scattering of gluons and photons is a one-loop process at LO. We decompose the scattering amplitude as

$$\mathcal{A}(1_g, 2_g, 3_g, 4_\gamma, 5_\gamma) = g_s e^2 \sum_{q=1}^{N_f} Q_q^2 f^{a_1 a_2 a_3} \sum_{\ell=1}^{\infty} \left(n_\epsilon \frac{\alpha_s}{4\pi} \right)^\ell A^{(\ell)}(1_g, 2_g, 3_g, 4_\gamma, 5_\gamma), \quad (4.2.4)$$

where

$$n_\epsilon = i \left(\frac{4\pi}{\mu_R} \right)^\epsilon e^{-\epsilon\gamma_E}, \quad (4.2.5)$$

with the renormalisation scale μ_R and the Euler–Mascheroni constant γ_E . The strong and EM vertex couplings are denoted by g_s and e respectively, N_f is the number of light quarks, Q_q is the electric charge of quarks of flavour q in units of the electron charge, a_i is the adjoint $SU(N_c)$ colour index of the i^{th} gluon, f^{a_i, a_j, a_k} are the $SU(N_c)$ structure constants, ℓ denotes the number of loops, and α_s is the strong coupling parameter. The one-loop diphoton amplitude can be obtained from permutations of pure-gluon scattering [251, 252].

We further expand the loop amplitudes in powers of N_c and N_f ,

$$\begin{aligned} A^{(1)}(1_g, 2_g, 3_g, 4_\gamma, 5_\gamma) &= A_1^{(1)}(1_g, 2_g, 3_g, 4_\gamma, 5_\gamma), \\ A^{(2)}(1_g, 2_g, 3_g, 4_\gamma, 5_\gamma) &= \underbrace{N_c A_1^{(2)}(1_g, 2_g, 3_g, 4_\gamma, 5_\gamma)}_{\text{LC}} \\ &\quad + \underbrace{\frac{1}{N_c} A_2^{(2)}(1_g, 2_g, 3_g, 4_\gamma, 5_\gamma) + N_f A_3^{(2)}(1_g, 2_g, 3_g, 4_\gamma, 5_\gamma)}_{\text{SLC}}. \end{aligned} \quad (4.2.6)$$

Surprisingly, the subleading-colour (SLC) two-loop amplitudes contain only planar integrals, while the LC contribution contains all of the four independent families shown in Fig. 4.1, including non-planar integrals. This pattern is the opposite to that of the quark-initiated channels computed in Refs. [208–210], for which the LC contributions involve only the planar integrals and are therefore simpler to compute. Providing a prediction for the gluon-initiated channel necessarily requires handling the most complicated integral families. A simple analysis of the colour factors of each of the three-gluon vertex diagrams shown in Fig. 4.2 illustrates how this pattern arises. Photons couple to any of the fermion propagators, and the colour factors remain the same. It can then be seen that non-planar contributions can come from the diagrams (a)–(c) only. Diagrams (d)–(e), which contribute to SLC, remain planar (allowing for permutations of the external momenta).

In our setup, we reduce directly to the FR where the UV and IR poles have been subtracted analytically. The poles take a particularly simple form since there is no tree-level process, the one-loop amplitudes are finite in ϵ , and the two-loop amplitudes are only divergent to $\mathcal{O}(\epsilon^{-2})$. The one- and two-loop FRs are given in




Figure 4.1: Independent integral families for the $gg \rightarrow g\gamma\gamma$ amplitude. The non-planar topologies (the second and fourth graphs) appear only in the LC amplitude.




Figure 4.2: The colour factor of each diagram in the $gg \rightarrow g\gamma\gamma$ follows from the representative three-gluon two-loop diagrams with a closed fermion loop shown here.

terms of the bare amplitudes [322–326] by

$$\begin{aligned} F^{(1)} &= A^{(1)}(1_g, 2_g, 3_g, 4_\gamma, 5_\gamma), \\ F^{(2)} &= A^{(2)}(1_g, 2_g, 3_g, 4_\gamma, 5_\gamma) - \left(I^{(1)} + \frac{3\beta_0}{2\epsilon} \right) A^{(1)}(1_g, 2_g, 3_g, 4_\gamma, 5_\gamma), \end{aligned} \quad (4.2.7)$$

where

$$\begin{aligned} \beta_0 &= \frac{11}{3}N_c - \frac{2}{3}N_f, \\ I^{(1)} &= -n_\Gamma(\epsilon) \left\{ \frac{N_c}{\epsilon^2} \left[\left(\frac{\mu_R^2}{-s_{12}} \right)^\epsilon + \left(\frac{\mu_R^2}{-s_{23}} \right)^\epsilon + \left(\frac{\mu_R^2}{-s_{13}} \right)^\epsilon \right] + 3\frac{\gamma_g}{\epsilon} \right\}, \end{aligned} \quad (4.2.8)$$

with

$$n_\Gamma(\epsilon) = \frac{e^{\epsilon\gamma_E}}{\Gamma(1-\epsilon)} = 1 - \frac{(\pi\epsilon)^2}{12} + \mathcal{O}(\epsilon^3), \quad \gamma_g = \frac{\beta_0}{2}, \quad (4.2.9)$$

where $\Gamma(z)$ is the gamma function, in the tHV scheme. The logarithms arising from the ϵ -expansion of $I^{(1)}$ can be analytically continued to the s_{12} channel by adding a small positive imaginary part to each s_{ij} . The β_0 term in the definition of the two-loop FR is defined by Eq. (1.1.18) and accounts for the strong coupling renormalisation. The FRs inherit from the amplitudes the decomposition in powers

of N_c and N_f given by Eq. (4.2.6),

$$\begin{aligned} F^{(1)}(1_g, 2_g, 3_g, 4_\gamma, 5_\gamma) &= F_1^{(1)}(1_g, 2_g, 3_g, 4_\gamma, 5_\gamma), \\ F^{(2)}(1_g, 2_g, 3_g, 4_\gamma, 5_\gamma) &= N_c F_1^{(2)}(1_g, 2_g, 3_g, 4_\gamma, 5_\gamma) \\ &\quad + \frac{1}{N_c} F_2^{(2)}(1_g, 2_g, 3_g, 4_\gamma, 5_\gamma) + N_f F_3^{(2)}(1_g, 2_g, 3_g, 4_\gamma, 5_\gamma). \end{aligned} \tag{4.2.10}$$

Our final results are presented in the tHV scheme, although we make the distinction between the dimension d of the loop integration and the spin dimension d_s arising from the numerator algebra (Section 1.3). We find that it is convenient to arrange terms by expanding the FR around the physical degrees of freedom of the gluon, $d_s = 2$. The one-loop and N_f two-loop FRs have only the $(d_s - 2)^0$ component. Thus, the expansions of all FRs are

$$\begin{aligned} F_1^{(1)} &= F_{1;0}^{(1)}, \\ F_k^{(2)} &= F_{k;0}^{(2)} + F_{k;1}^{(2)} (d_s - 2) \quad \forall k \in \{1, 2\}, \\ F_3^{(2)} &= F_{3;0}^{(2)}. \end{aligned} \tag{4.2.11}$$

4.3 Computational setup and amplitude reduction

We take a diagrammatic approach to the calculation of the amplitude along the lines of previous work [327, 328]. Here we briefly summarise the steps, which are sketched in Fig. 4.3, and refer the reader to Ref. [328] for further details. All Feynman diagrams are generated using `QGRAF` [329] and subsequently processed using a combination of in-house `Mathematica` and `FORM` [330, 331] scripts. In total, including contributions from ghost diagrams, we find 50 diagrams at one loop and 1527 at two loops. Aided by the `Spinney` [332] package to perform the 't Hooft algebra, the numerators are written for each independent helicity configuration. From the loop denominator structure we assign an integral topology to each diagram. At this point, the diagram numerators are linear combinations of monomials in loop-momentum dependent scalar and spinor products with coefficients depending only on external momenta. These coefficients are loaded into a dataflow graph using `FiniteFlow` [167]. This enables numerical sampling over FFs, thus sidestepping analytically complicated intermediate expressions in further steps. We rewrite loop-momentum dependent monomials into inverse propagator denominators and a choice of ISPs. The required mapping of the coefficients is performed numerically within the dataflow framework. After summing all diagrams and dropping scaleless integrals,

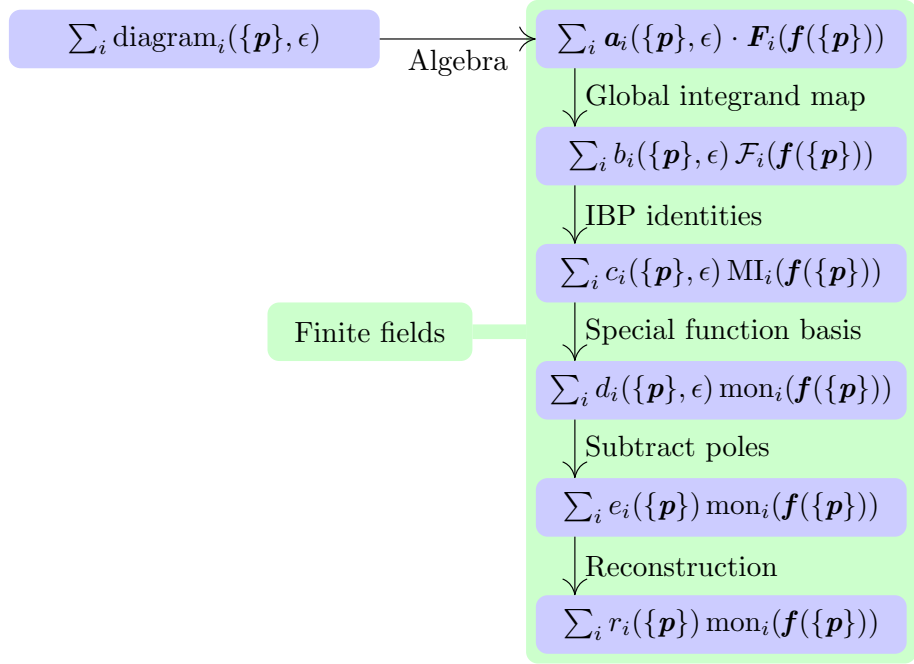


Figure 4.3: Schematic flow chart of the steps in our assembly of the FRs. We start by summing over all Feynman diagrams, then algebraically manipulating to a linear combination of tensor integrals \mathbf{F}_i , which are composed of special functions $\mathbf{f}(\{\mathbf{p}\})$. The coefficients a_i are rational functions of the kinematics $\{\mathbf{p}\}$ and the dimensional regulator ϵ . All computation from this point on is performed in FFs. We reduce to scalar integrals $\mathcal{F}_i(\{\mathbf{p}\})$, then to a set of LI MIs using IBP identities. We project these onto a special function basis (the pentagon functions) to obtain a sum over LI monomials of the special functions $\text{mon}_i(\mathbf{f}(\{\mathbf{p}\}))$. We perform the IR subtraction to recover the tHV FRs. Finally, we reconstruct the coefficients to obtain compact analytic expressions $r_i(\{\mathbf{p}\})$ for them.

we arrive at an expression ready for IBP reduction.

The reduction to MIs was obtained using an improved version of the Laporta algorithm [163]. For most integral families, we generated identities containing no higher power of propagators with respect to those appearing in the amplitude, following ideas proposed in Refs. [187, 288, 290]. These identities were found using the Baikov representation of loop integrals, for which identities (i) without higher powers of propagators and (ii) without dimension-shifted integrals can be found by solving polynomial equations called syzygy equations. Closed-form solutions to both of these constraints are separately known. Indeed, the solution of (i) is straightforward and the solution for (ii) has been found in Ref. [291]. The two syzygy solutions need to be combined for generating identities that satisfy both constraints. For this purpose, we used a custom syzygy solver that implements the algorithm in Ref. [289] using `FiniteFlow` [167]. More details on this method can be found in Refs. [187, 288, 290, 291]. The application of the syzygy technique leads to a substantial reduction in the size of the IBP system in the planar sector, which improves both the speed of solving the system and memory usage. This improved performance in the planar sectors by almost a factor of ten, which was sufficient for the current calculation, without the need to extend it to the non-planar families—though we expect this will be necessary for future applications.

For each integral family, we generated integral identities for only one permutation of the external legs. Numerical solutions for all the permutations contributing to an amplitude were found by solving the systems of equations several times, with different numerical inputs for the invariants. Mappings between MIs with different permutations of external legs are applied afterwards to obtain a result in terms of a minimal set of MIs.

For each phase-space point evaluated on an FF, we reconstruct the full dependence on the dimensional regulator ϵ of the amplitude reduced onto MIs before substituting their expressions in terms of the basis of special functions and computing the Laurent expansion in ϵ . With this setup, fewer numerical solutions of the IBP identities are needed in order to reconstruct analytic results for the amplitude. This is due to the fact that if we were to expand the integrals into the pentagon functions before performing the Laurent expansion in ϵ of the final coefficients, it would complicate the dependence on ϵ of the result in this intermediate stage.

To make use of the FF arithmetic we must have a rational parametrisation of the external kinematics. As discussed in Section 1.5.7, we parametrise the kinematics

using MTVs,

$$\begin{aligned}
s_{12} &= x_1, \\
s_{23} &= x_1 x_4, \\
s_{34} &= \frac{x_1}{x_2} (x_4 + x_3 x_4 + x_2 x_3 x_5 - x_2 x_3), \\
s_{45} &= x_1 x_5, \\
s_{15} &= x_1 x_3 (x_2 - x_4 + x_5), \\
\text{tr}_5 &= -\frac{x_1^2}{x_2} [x_2 x_4 (1 + 2x_3) - x_4 (1 + x_3) (x_4 - x_5) + x_2^2 x_3 (-1 + x_5)].
\end{aligned} \tag{4.3.1}$$

We stress that the pseudoscalar invariant tr_5 , and hence the square root of the Gram determinant Δ (see Section 1.5.1), is a rational function of the x_i variables. Moreover, since x_1 is the only dimensionful variable, we can set it equal to unity and recover the dependence on it after the reconstruction by dimensional analysis. Further details on the momentum twistor parametrisation are presented in Section 4.4. In the following sections, we will consider all coefficients of the special functions to be rational functions of the variables x_i .

4.4 Momentum twistor parametrisation

MTVs were introduced in Section 1.5.7. This construction follows Refs. [130, 131, 308]. It is possible to fix all but five of the entries of the momentum twistor matrix, Eq. (1.5.70), at five-point. Explicitly, we choose the form,

$$Z = \begin{pmatrix} \lambda_i \\ \mu_i \end{pmatrix}_{i \in \{1, \dots, 5\}} = \begin{pmatrix} 1 & 0 & \frac{1}{x_1} & \frac{1+x_2}{x_1 x_2} & \frac{1+x_3(1+x_2)}{x_1 x_2 x_3} \\ 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & \frac{x_4}{x_2} & 1 \\ 0 & 0 & 1 & 1 & \frac{x_4 - x_5}{x_4} \end{pmatrix}. \tag{4.4.1}$$

The parametrisation used in this work has some benefits: the only dimensionful quantity is x_1 (recall that we have set this to unity) and all holomorphic quantities are described using only x_1, x_2, x_3 . For real kinematics, only x_2 and x_3 are complex, while the other three are real. Notice that the conversion between the MTVs and spinor-helicity expressions, Eq. (4.3.1), is only invertible for phase-free quantities as the phase information is lost when translating to MTVs. Thus, for an invertible

mapping, we may use the following relations,

$$\begin{aligned}
x_1 &= s_{12}, \\
x_2 &= -\frac{\text{tr}_+(p_1, p_2, p_3, p_4)}{s_{12}s_{34}}, \\
x_3 &= -\frac{\text{tr}_+(p_1, p_3, p_4, p_5)}{s_{13}s_{45}}, \\
x_4 &= \frac{s_{23}}{s_{12}}, \\
x_5 &= \frac{s_{45}}{s_{12}},
\end{aligned} \tag{4.4.2}$$

where $\text{tr}_+(p_i, p_j, p_k, p_l)$ is defined in Eq. (1.5.17).

In our work, we express the helicity amplitudes in terms of MTVs x_i . The phase information can be restored by multiplying and dividing by a suitable phase factor,

$$\mathcal{A} = \Phi(\lambda_i, \tilde{\lambda}_i) \left(\frac{\mathcal{A}(x_i)}{\Phi(x_i)} \right), \tag{4.4.3}$$

where \mathcal{A} is a helicity amplitude—or in general some object with a non-trivial phase—and Φ is an arbitrary factor with the same helicity weights (defined in Section 1.5.4) as \mathcal{A} . The quantities $\mathcal{A}(x_i)$ and $\Phi(x_i)$ are both written in terms of MTVs. Their ratio is phase-free and can thus be expressed in terms of the scalar and pseudoscalar invariants, s_{ij} and tr_5 , for example through Eq. (4.4.2), or evaluated directly in MTVs. The phase $\Phi(\lambda_i, \tilde{\lambda}_i)$ can be constructed using Eq. (1.5.58), is written in terms of the spinor-helicity variables, and carries all the phase information of \mathcal{A} . The phase $\Phi(x_i)$ can be obtained from $\Phi(\lambda_i, \tilde{\lambda}_i)$ by applying the MTV parametrisation, which in our case is given by the mapping Eq. (4.3.1).

For the aid of comparisons with the data presented in this chapter, the specific choices of the amplitude phases $\Phi(\lambda_i, \tilde{\lambda}_i)$ were

$$\begin{aligned}
\Phi(1_g^+, 2_g^+, 3_g^+, 4_\gamma^+, 5_\gamma^+) &= \frac{1}{\langle 12 \rangle \langle 23 \rangle \langle 34 \rangle \langle 45 \rangle \langle 51 \rangle}, \\
\Phi(1_g^-, 2_g^+, 3_g^+, 4_\gamma^+, 5_\gamma^+) &= \frac{[23] \langle 13 \rangle}{\langle 23 \rangle \langle 34 \rangle \langle 45 \rangle \langle 15 \rangle}, \\
\Phi(1_g^+, 2_g^+, 3_g^+, 4_\gamma^-, 5_\gamma^+) &= \frac{[15] \langle 14 \rangle}{\langle 12 \rangle \langle 23 \rangle \langle 34 \rangle \langle 15 \rangle}, \\
\Phi(1_g^-, 2_g^-, 3_g^+, 4_\gamma^+, 5_\gamma^+) &= \frac{\langle 12 \rangle^3}{\langle 23 \rangle \langle 34 \rangle \langle 45 \rangle \langle 51 \rangle}, \\
\Phi(1_g^+, 2_g^+, 3_g^+, 4_\gamma^-, 5_\gamma^-) &= \frac{\langle 45 \rangle^3}{\langle 12 \rangle \langle 23 \rangle \langle 34 \rangle \langle 51 \rangle}, \\
\Phi(1_g^-, 2_g^+, 3_g^+, 4_\gamma^-, 5_\gamma^+) &= \frac{\langle 14 \rangle^4}{\langle 12 \rangle \langle 23 \rangle \langle 34 \rangle \langle 45 \rangle \langle 51 \rangle}.
\end{aligned} \tag{4.4.4}$$

The MHV trees were used along with contrived expressions for the UHVs.

4.5 Analytic reconstruction over finite fields

In this section, we present three general strategies to optimise the reconstruction over FFs of the rational coefficients in the FRs. At this stage, each component $F(x)$ of the two-loop FR is expressed as

$$F(x) = \sum_i r_i(x) \text{mon}_i(f), \quad (4.5.1)$$

where r_i are rational functions of the MTVs x , and $\text{mon}_i(f)$ are LI monomials of the pentagon functions. The entire chain of operations is implemented over FFs in the framework `FiniteFlow`. We therefore have a numerical algorithm which evaluates the rational coefficients $r_i(x)$ modulo some prime number. The final step consists of reconstructing the analytic expression of the rational coefficients from a sufficient number of numerical evaluations. We employ `FiniteFlow`'s multi-variate functional reconstruction algorithms, supplemented with three strategies to reduce the number of required sample points: we determine the linear relations among the rational coefficients and an ansatz, use univariate slices to identify the factors belonging to another ansatz, and perform a univariate PFD on the fly. In the following subsections we discuss thoroughly each of these procedures and their application to two-loop diphoton FRs.

4.5.1 Linear relations among the rational coefficients

The representation of the FRs in terms of rational coefficients and special function monomials given by Eq. (4.5.1) is not optimal. The special function monomials do not all appear independently. They are present only in a number of LI combinations that is typically much smaller than the total number of monomials. As a result, the rational coefficients r_i in the FRs are not LI. Expressing the FRs in terms of a set of LI rational coefficients not only leads to more compact expressions, but may also simplify their reconstruction.

We can determine the linear relations among the rational coefficients $\{r_i(x)\}$ of the special function monomials by solving a linear fit problem,

$$\sum_i a_i r_i(x) = 0. \quad (4.5.2)$$

Since the coefficients of the linear relations a_i are rational numbers, they require substantially fewer sample points to be reconstructed with respect to the rational

coefficients themselves. We can then use these relations to express the rational coefficients in terms of a set of LI ones, which remain to be reconstructed. Choosing the latter to be the simplest—*i.e.* those with the lowest polynomial degrees—may reduce the number of sample points required for the reconstruction.

This strategy can be further refined by supplying an ansatz for the rational coefficients. We then fit the linear relations among the rational coefficients of the FRs and the coefficients of the ansatz, which we denote by $\{e_j(x)\}$, as

$$\sum_i a_i r_i(x) + \sum_j b_j e_j(x) = 0, \quad (4.5.3)$$

with $a_i, b_j \in \mathbb{Q}$. In the best case scenario, all the rational coefficients r_i can be expressed in terms of the ansatz coefficients e_j and no further reconstruction needs to be performed. Even when the ansatz does not entirely cover the rational coefficients, it may still lower the degrees of the LI coefficients which have to be reconstructed. The ansatz can be constructed from the tree-level amplitude and the rational coefficients of the one-loop amplitudes up to $\mathcal{O}(\epsilon^2)$ from the analysis of the leading singularities [151, 333–335] or from other related amplitudes. In the diphoton case, we can use the two-loop five-gluon amplitudes. At one loop, the $3g2\gamma$ amplitudes can be expressed in terms of permutations of the five-gluon ones [251, 252]. While this is no longer true at two loops, we find there is an important overlap between the rational coefficients of the $3g2\gamma$ amplitudes and those of the five-gluon ones. We use the rational coefficients of the LC two-loop five-gluon amplitudes as ansätze in the linear relations; all two-loop five-parton amplitudes are available analytically at LC [312–314, 316, 336–339] and we made use of the independent results discussed in Chapter 6.

4.5.2 Matching factors on univariate slices

The pole structure of the pentagon functions is determined by the letters of the pentagon alphabet [279]. The pentagon functions (or their discontinuities) may in fact have logarithmic singularities in the phase-space points where one of the letters vanishes. For this reason, it is natural to expect that the poles of the rational coefficients should be similarly linked to the pentagon alphabet. Indeed, we observe that the denominators of the rational coefficients in front of the pentagon functions factorise into a product of letters of the pentagon alphabet. In other words, each rational coefficient $r(x)$ has the form,

$$r(x) = \frac{n(x)}{\prod_k \ell_k^{e_k(x)}}, \quad (4.5.4)$$

where e_k are integers, $n(x)$ is a polynomial in the variables x , and $\{\ell_k\}$ is an ansatz of factors from the pentagon alphabet. The exponents e_k in Eq. (4.5.4) may in general be negative, corresponding to factors in the numerator. We use the following ansatz for the factors¹,

$$\begin{aligned} \{\ell_k(x)\} = \{ & \langle 1\ 2 \rangle, \langle 1\ 3 \rangle, \langle 1\ 4 \rangle, \langle 1\ 5 \rangle, \langle 2\ 3 \rangle, \langle 2\ 4 \rangle, \langle 2\ 5 \rangle, \langle 3\ 4 \rangle, \langle 3\ 5 \rangle, \langle 4\ 5 \rangle, [1\ 2], [1\ 3], [1\ 4], \\ & [1\ 5], [2\ 3], [2\ 4], [2\ 5], [3\ 4], [3\ 5], [4\ 5], s_{12} - s_{34}, s_{12} - s_{35}, s_{12} - s_{45}, \\ & s_{13} - s_{24}, s_{13} - s_{25}, s_{13} - s_{45}, s_{14} - s_{23}, s_{14} - s_{25}, s_{14} - s_{35}, s_{15} - s_{23}, \\ & s_{15} - s_{24}, s_{15} - s_{34}, s_{23} - s_{45}, s_{24} - s_{35}, s_{25} - s_{34}, \text{tr}_5 \}. \end{aligned} \quad (4.5.5)$$

The exponents e_k in the ansatz Eq. (4.5.4) can be determined by reconstructing $r(x)$ on a univariate slice modulo some prime number [314]. The univariate slice is defined by parametrising the variables in terms of a single parameter t ,

$$\{x_i(t) = a_i + b_i t\}, \quad (4.5.6)$$

for constant a_i and b_i . The constants are chosen randomly in the FF to avoid artificial simplifications. The dependence on t is chosen to be linear so that the degrees of the numerator and denominator of

$$r(t) := r(x(t)) \quad (4.5.7)$$

correspond to the total degrees of r in x . Matching the reconstructed $r(t)$ with the ansatz Eq. (4.5.4) evaluated on the same slice allows to determine the exponents e_k straightforwardly. With a univariate reconstruction on just one prime field we can thus infer a lot of information about the analytic form of the rational coefficients: the denominators are entirely fixed, and typically some factors of the numerators are determined as well. What remains to be reconstructed therefore requires fewer sample points.

4.5.3 Univariate partial fraction decomposition over finite fields

PFD is a standard and powerful tool for the simplification of rational functions. The PFD is however not unique in the multivariate case. Its application to the multivariate rational functions in scattering amplitudes is therefore not straightforward. The necessity to simplify the rational coefficients of two-loop five-particle scatter-

¹Note that the ansatz in Eq. (4.5.5) is independent of the rational parametrisation of the kinematics, Eq. (4.3.1). The list of independent polynomials used in the factor matching on the univariate slice is generated by evaluating this list using the specific parametrisation.

ing amplitudes has recently spurred several approaches to handle the multivariate case efficiently [293, 315, 340], based upon Leinartas' algorithm [341, 342]. These algorithms rely on algebraic geometry techniques, such as multivariate polynomial division and Gröbner bases, and require an arbitrary choice of a monomial ordering.

Our main goal in this work is to simplify the reconstruction of the rational coefficients over FFs. In other words, we want to reconstruct the rational coefficients on the fly, directly in a form which is decomposed in partial fractions. The simplification of the resulting analytic expressions comes as a welcome by-product. We observe that a univariate PFD is sufficient for this purpose. The advantage is that it can be straightforwardly implemented over FFs, avoiding all algebraic geometry complications. The only arbitrary choice that remains to be made is to choose which variable we will partial fraction with respect to. This can be chosen by observing the impact of the PFD with respect to each variable separately on the lower order amplitudes. With the parametrisation of the kinematics in terms of momentum twistors, Eq. (4.4.1), we find it most convenient to decompose in partial fractions with respect to x_4 .

We now discuss our algorithm to reconstruct the univariate PFD of a multivariate rational function r from its numerical evaluations over FFs. The algorithm requires as input an ansatz for the factors which may appear in the denominator of r . Only those factors which depend on the variable with respect to which the PFD is being performed are strictly necessary. Informed guesses of other factors may further simplify the reconstruction. In the application to massless two-loop five-particle scattering amplitudes, the factor ansatz can be inferred from the letters of the pentagon alphabet [279]. We use the factors in Eq. (4.5.5).

Let r be a rational function of the set of variables,

$$x = \{x_i\}_{i=1}^n . \quad (4.5.8)$$

In this work, the x_i are the MTVs defined by Eq. (4.3.1) with $n = 5$, but we outline the algorithm in general. The goal is to decompose r in partial fractions with respect to one of the variables, say $y := x_k$. We denote the remaining variables by

$$\bar{x} = \{x_i\}_{i=1}^n \setminus \{y\} . \quad (4.5.9)$$

We may not know the analytic expression of r , but we must be able to evaluate it numerically modulo some prime number. Let

$$\{\ell_i(\bar{x}, y)\}_{i=1}^m \quad (4.5.10)$$

be an ansatz for the m factors which may appear in the denominator of r . Without loss of generality, we assume that the ℓ_i are irreducible polynomials over \mathbb{Q} . In other

words, we assume that r has the form

$$r(\bar{x}, y) = \frac{N(\bar{x}, y)}{\prod_{i=1}^m \ell_i^{e_i}(\bar{x}, y)}, \quad (4.5.11)$$

where the exponents $e_i \in \mathbb{Z}$, and $N(\bar{x}, y)$ is a function which depends polynomially on y and rationally on \bar{x} . The ansatz, Eq. (4.5.10), may catch some of the factors in the numerator of $r(\bar{x}, y)$, corresponding to negative values of e_i . This lowers the total degree of $N(\bar{x}, y)$ and eventually simplifies its reconstruction, but is not necessary for the PFD with respect to y . Similarly, the ansatz may cover all the factors in the denominator of r , so that $N(\bar{x}, y)$ is a polynomial in \bar{x} as well as y . What is necessary for the PFD algorithm to work is that the ansatz contains all the factors in the denominator of r which depend on y . We denote this subset by

$$\Lambda_y = \{i \in \{1, \dots, m\} : e_i > 0 \wedge \deg_y[\ell_i(\bar{x}, y)] > 0\}, \quad (4.5.12)$$

where $\deg_y[\ell_i(\bar{x}, y)]$ is the degree in y of the polynomial $\ell_i(\bar{x}, y)$.

The first step consists of fixing the exponents e_i in the ansatz Eq. (4.5.11). We do this through the procedure discussed in Section 4.5.2. In the second step, we determine the degree in y of $N(\bar{x}, y)$ in the ansatz Eq. (4.5.11),

$$d_N := \deg_y[N(\bar{x}, y)]. \quad (4.5.13)$$

We recall that $N(\bar{x}, y)$ is by construction a polynomial in y . We compute its degree in y by reconstructing it on another univariate slice, this time where only y varies,

$$\{\bar{x}_i(t) = \bar{a}_i, \quad y(t) = t\}, \quad (4.5.14)$$

with a_i chosen randomly from the FF. Clearly,

$$d_N = \deg_t[N(t)], \quad N(t) := N(\bar{x} = \bar{a}, y = t). \quad (4.5.15)$$

We introduce the short-hand notation,

$$d_i := \deg_y[\ell_i(\bar{x}, y)], \quad d_{\Lambda_y} := \sum_{i \in \Lambda_y} e_i d_i, \quad (4.5.16)$$

for the degrees of the denominator factors $\ell_i(\bar{x}, y)$ in y .

Using the information about the factors in the denominator of r and the degree in y of its numerator, we construct the following ansatz for the PFD of r with respect to y ,

$$r(\bar{x}, y) = \sum_{i \in \Lambda_y} \sum_{j=1}^{e_i} \sum_{k=0}^{d_i-1} \frac{U_{ijk}(\bar{x}) y^k}{\ell_i^j(\bar{x}, y)} + R(\bar{x}) + \sum_{h=1}^{d_N-d_{\Lambda_y}} V_h(\bar{x}) y^h, \quad (4.5.17)$$

where $U_{ijk}(\bar{x})$, $R(\bar{x})$ and $V_h(\bar{x})$ are unknown rational functions of \bar{x} . The rightmost term in Eq. (4.5.17) is required only if $d_N > d_{\Lambda_y}$, *i.e.* only if the numerator of r has a higher degree in y than the denominator.

The last step of the algorithm consists of reconstructing the analytic dependence on \bar{x} of the unknown coefficients in the ansatz, Eq. (4.5.17), from the numerical evaluations of $r(\bar{x}, y)$. To solve this linear fit problem, we use the algorithm implemented in the `FiniteFlow` framework [167]. The solution comes in the form of an algorithm which numerically evaluates $U_{ijk}(\bar{x})$, $R(\bar{x})$ and $V_h(\bar{x})$. The rational reconstruction may be simplified by first reconstructing the coefficients on a univariate slice where all the remaining variables \bar{x} vary, and using that to match them with those factors in the ansatz, Eq. (4.5.10), which depend only on \bar{x} . This may lower the total degree of the functions that need to be reconstructed.

In addition to the factors in the original ansatz, Eq. (4.5.10), the coefficients of the PFD, Eq. (4.5.17), may also contain spurious factors (Section 2.1.2). For instance, consider the toy example,

$$\frac{1}{(y-a)(y-b)} = \frac{1}{(a-b)(y-a)} - \frac{1}{(a-b)(y-b)}, \quad (4.5.18)$$

where a and b are arbitrary constants such that $a \neq b$. In this example, the inspection of the left-hand side indicates $\{y-a, y-b\}$ as ansatz for the irreducible denominator factors. The PFD however contains a factor of $a-b$ in the denominator, which arises from the residue of the function at the root of either of the denominator factors. Clearly $a=b$ is a spurious singularity, manifestly absent on the left-hand side and produced by the PFD. In general, we can determine the potential spurious factors by evaluating the factors in the ansatz, Eq. (4.5.10), which depend on y at their zeros,

$$\{\ell_i(\bar{x}, y_k^*)\}_{i \in \Lambda_y, k \in \Lambda_y^1, i \neq k}, \quad (4.5.19)$$

where y_k^* is the zero of $\ell_k(\bar{x}, y)$,

$$\ell_k(\bar{x}, y_k^*) = 0, \quad (4.5.20)$$

and Λ_y^1 is the subset of factors which depend linearly on y ,

$$\Lambda_y^1 = \{i \in \Lambda_y : \deg_y[\ell_i(\bar{x}, y)] = 1\}. \quad (4.5.21)$$

The restriction to zeros of linear functions of y is due to the facts that the ℓ_i are irreducible polynomials over \mathbb{Q} and that we are factoring over \mathbb{Q} . The zeros of higher-degree irreducible polynomials would introduce algebraic and/or complex dependence.

In practice, we observe that determining the spurious factors does not simplify the reconstruction. The greatest part of the denominators of the coefficients in the PFD, Eq. (4.5.17), is in fact determined by the original ansatz, Eq. (4.5.10). What remains after they are multiplied away has a total degree which is typically lower than that of the numerators, which therefore dominates the determination of the number of sample points required for the reconstruction. While it is possible to determine entirely the denominators of the coefficients in Eq. (4.5.17), it would not reduce the number of required sample points substantially, and for this reason we refrain from doing so.

Having determined as many factors as possible in the coefficients of the PFD, we multiply them away and reconstruct the remainder using the multivariate rational reconstruction algorithms implemented in `FiniteFlow`. It is important to stress that the algorithm which evaluates the coefficients of the PFD contains the solution of a linear fit. For each numerical value of \bar{x} , Eq. (4.5.17) is sampled for several numerical values of y , roughly as many times as the number of unknowns. This generates a linear system of equations for the unknowns evaluated at the chosen value of \bar{x} . The redundant equations are removed after the learning phase. Because it requires several evaluations of the original functions, the reconstruction on the univariate slices in the intermediate steps of the algorithm has a higher computational cost than directly evaluating r . On the other hand, the coefficients of the PFD depend on one fewer variable than the original function r , and may have substantially lower degrees. As a result of all these aspects, the PFD may be outperformed by a direct reconstruction for simple functions, but becomes increasingly convenient as the complexity of the functions rises. It is well suited for application at two loops.

4.5.4 Summary and impact of the reconstruction strategy

The techniques discussed in the previous sections are general and can be applied to any rational reconstruction problem, in combination or separately. In order to reconstruct the rational coefficients of the two-loop diphoton FRs we apply them consecutively as follows.

Stage 1. We fit the linear relations among the rational coefficients with an ansatz, as discussed in Section 4.5.1. We begin with the $(d_s - 2)^1$ components and use the coefficients of the two-loop LC five-gluon FRs as ansätze. For the $(d_s - 2)^0$ components, which are more complicated, we add to the ansatz the $(d_s - 2)^1$ -coefficients already reconstructed.

Stage 2. We guess the factors from the ansatz Eq. (4.5.5) by reconstructing a univariate slice (Section 4.5.2) and multiply them away.

Finite remainder	Original	Stage 1	Stage 2	Stage 3*	Stage 4*
$F_{1;1}^{(2)}(1_g^-, 2_g^-, 3_g^+, 4_\gamma^+, 5_\gamma^+)$	69/60	28/20	24/0	19/10	11/5
$F_{1;0}^{(2)}(1_g^-, 2_g^-, 3_g^+, 4_\gamma^+, 5_\gamma^+)$	78/69	44/35	43/0	21/10	16/9
$F_{1;1}^{(2)}(1_g^-, 2_g^+, 3_g^+, 4_\gamma^-, 5_\gamma^+)$	59/55	30/27	29/0	18/15	17/4
$F_{1;0}^{(2)}(1_g^-, 2_g^+, 3_g^+, 4_\gamma^-, 5_\gamma^+)$	89/86	38/36	38/0	20/16	17/3
$F_{1;1}^{(2)}(1_g^+, 2_g^+, 3_g^+, 4_\gamma^-, 5_\gamma^-)$	40/42	25/27	25/0	15/18	15/0
$F_{1;0}^{(2)}(1_g^+, 2_g^+, 3_g^+, 4_\gamma^-, 5_\gamma^-)$	66/66	32/33	32/0	13/13	12/3

Table 4.1: Maximal numerator/denominator polynomial degrees of the rational coefficients of the most complicated FRs at each stage of our reconstruction strategy. The column “original” refers to the rational coefficients prior to any optimisation. The asterisk highlights that, after the PFD in Stage 3, the coefficients to be reconstructed depend on one fewer variable.

Stage 3. We partial fraction on the fly with respect to x_4 , applying the algorithm presented in Section 4.5.3. The coefficients to be reconstructed after this stage are those in the ansatz for the PFD Eq. (4.5.17), and depend on one fewer variable.

Stage 4. We reconstruct another univariate slice and perform an additional factor guessing, as in Stage 2.

The drop in the complexity of the rational coefficients after each stage for the most complicated two-loop diphoton FRs, which are in the MHV configurations, is illustrated in Table 4.1. As proxy for the complexity of the coefficients we use the maximal numerator/denominator polynomial degrees, which can be evaluated by reconstructing univariate slices as discussed in Section 4.5.3.

Interestingly, we observe that the coefficients of the SLC $3g2\gamma$ two-loop FRs $F_2^{(2)}$ can be expressed in terms of those of the LC two-loop five-gluon FRs. The coefficients of the LC $3g2\gamma$ two-loop remainders $F_1^{(2)}$ instead are not entirely fixed by the five-gluon ones, but using the latter as ansätze in the linear relations reduces significantly the maximal polynomial degrees of the coefficients which remain to be reconstructed.

As can be appreciated in Table 4.1, our strategy leads to a substantial drop in the polynomial degrees. Furthermore, the coefficients to be reconstructed after the PFD (Stage 3) depend on one fewer variable. This makes the decrease in the number of sample points required for the reconstruction even more pronounced. The price to pay for this is that performing the PFD increases the evaluation time per point, as discussed at the end of Section 4.5.3. With our setup we observe that, for the most complicated FRs, the evaluation times grows roughly by one order of magnitude, while the number of sample points required for the reconstruction decreases by two

orders of magnitude. This leads to an overall gain of roughly one order of magnitude in the reconstruction time¹. We stress that the evaluation time relevant here is that of the algorithm which evaluates the rational coefficients over FFs, not the final evaluation time of the FRs. Once the reconstruction is completed, in fact, the rational coefficients are evaluated from their analytic expressions. We will discuss the evaluation time of the FRs in Section 4.7.

Our approach therefore leads to an important simplification in the reconstruction of the rational coefficients. Moreover, the ensuing analytic expressions are dramatically more compact. This makes them suitable for compilation in a C++ library, an essential step for their phenomenological application, which we discuss in Chapter 5.

4.6 Compact analytic expressions for the all-plus configuration

Prior to discussing the numerical implementation of all two-loop helicity amplitudes, we would like to comment on the all-plus amplitude, which displays a particularly simple analytic form. We find that the structures appearing are closely related to those appearing in the five-gluon all-plus amplitudes at one [343–346] and two loops [312, 336–338]. We present the FRs in the expansion around $d_s = 2$ (Section 4.2).

The all-plus amplitude is finite and rational at one loop. The FR can be written as

$$F_{1;0}^{(1)}(1_g^+, 2_g^+, 3_g^+, 4_\gamma^+, 5_\gamma^+) = -2 \frac{[45]^2}{\langle 12 \rangle \langle 23 \rangle \langle 31 \rangle}. \quad (4.6.1)$$

Remarkably, this amplitude is invariant under conformal transformations, and the expression given here exhibits this property in a manifest way [346]. If all masses are neglected, the SM Lagrangian is conformally invariant [347]. This symmetry is obscured at loop level by the appearance of scales associated with the divergences and it is therefore rather surprising to observe it in a one-loop amplitude. One might naïvely suppose that this is a consequence of the finiteness of the all-plus one-loop amplitudes. Yet, the single-minus one-loop amplitudes are equally finite, but they are not conformally invariant. This phenomenon still calls for an explanation.

¹To give a sense of the absolute scale of the improvement, we quote the explicit number of sample points required for the $- + + - +$ helicity configuration for the LC $(d_s - 2)^0$ component, $F_{1;0}^{(2)}$. Reconstructing after Stage 2 would have required 57 291 sample points in four variables while reconstructing after Stage 4 requires 518 sample points in three variables. Note that each of the 518 points requires a univariate fit in the additional variable and so the overall improvement is around a factor of 10.

These properties are discussed in detail in Ref. [346], where the authors prove that the n -gluon all-plus amplitudes in QCD are conformally invariant at one loop. Since the diphoton amplitudes can be expressed as permutations of pure-gluon scattering [251, 252] and the conformal generators commute with permutations, all considerations regarding conformal symmetry extend to the diphoton case.

At two-loop order, the LC $(d_s - 2)^0$ contribution is the only one involving transcendental functions (Section 1.6). Its expression is remarkably simple,

$$F_{1;0}^{(2)}(1_g^+, 2_g^+, 3_g^+, 4_\gamma^+, 5_\gamma^+) = \frac{[45]^2}{\langle 12 \rangle \langle 23 \rangle \langle 31 \rangle} \sum_{\text{cyclic}(123)} F_{\text{box}}(s_{12}, s_{23}; s_{45}), \quad (4.6.2)$$

where the sum runs over the cyclic permutations of $(1, 2, 3)$, and

$$F_{\text{box}}(s_{12}, s_{23}; s_{45}) = \text{Li}_2\left(1 - \frac{s_{12}}{s_{45}}\right) + \text{Li}_2\left(1 - \frac{s_{23}}{s_{45}}\right) + \ln^2\left(\frac{s_{12}}{s_{23}}\right) + \frac{\pi^2}{6} \quad (4.6.3)$$

is the finite part of the one-loop box with an off-shell leg. The dilogarithm is defined in Eq. (1.6.3). The analytic continuation of the box functions to any scattering region can be achieved by adding a small positive imaginary part to each two-particle momentum invariant,

$$s_{ij} \rightarrow s_{ij} + i0^+. \quad (4.6.4)$$

The other partial amplitudes at two loops are rational,

$$\begin{aligned} F_{1;1}^{(2)}(1_g^+, 2_g^+, 3_g^+, 4_\gamma^+, 5_\gamma^+) &= -\frac{[45]^2}{\langle 12 \rangle \langle 23 \rangle \langle 31 \rangle} - \frac{1}{2} F_{3;0}^{(2)}(1_g^+, 2_g^+, 3_g^+, 4_\gamma^+, 5_\gamma^+), \\ F_{2;0}^{(2)}(1_g^+, 2_g^+, 3_g^+, 4_\gamma^+, 5_\gamma^+) &= 0, \\ F_{2;1}^{(2)}(1_g^+, 2_g^+, 3_g^+, 4_\gamma^+, 5_\gamma^+) &= -3 \frac{[45]^2}{\langle 12 \rangle \langle 23 \rangle \langle 31 \rangle} - \frac{1}{2} \frac{\text{tr}_5(p_1, p_2, p_3, p_4 - p_5) \langle 45 \rangle}{\langle 14 \rangle \langle 15 \rangle \langle 24 \rangle \langle 25 \rangle \langle 34 \rangle \langle 35 \rangle}, \\ F_{3;0}^{(2)}(1_g^+, 2_g^+, 3_g^+, 4_\gamma^+, 5_\gamma^+) &= \frac{1}{3} \text{tr}_5(p_1, p_2, p_3, p_4 - p_5) \sum_{\text{cyclic}(123)} \frac{1}{\langle 23 \rangle^2 \langle 14 \rangle \langle 15 \rangle \langle 45 \rangle}, \end{aligned} \quad (4.6.5)$$

where $\text{tr}_5(p_i, p_j, p_k, p_l)$ is defined in Eq. (1.5.15). The peculiar simplicity of this amplitude at two loops follows from the fact that it vanishes at tree level and it is rational in four dimensions at one loop. The one-loop amplitude can in fact be used as an effective on-shell vertex in four-dimensional unitarity [148, 149, 336]. In this way, the cuts of the two-loop amplitude become one-loop cuts with an insertion of the effective vertex. The one- and two-loop all-plus FRs are thus treated as tree-level and one-loop objects, respectively. As a result, the special functions appearing in the FR at two loops can have at most transcendental weight two (up to $\mathcal{O}(\epsilon^0)$).

Moreover, the rational coefficients of the transcendental functions can be shown through four-dimensional unitarity to be given by (permutations of) the one-loop all-plus FR. Thus, they inherit the symmetry under conformal transformations from the one-loop amplitude. These properties are manifest in our explicit expressions, Eqs. (4.6.1) and (4.6.2). Complementing four-dimensional unitarity with recursion relations for the rational terms allows us to compute the two-loop all-plus FRs in the purely gluonic case avoiding altogether the computation of the two-loop integrals [336, 338]. Some results are available even for amplitudes involving more than five positive-helicity gluons [348–353].

Amplitudes with a single minus helicity share some of the simplicity of the all-plus case. They also vanish at tree level, and are finite and rational at one loop. As a result, they also have maximum transcendental weight two at two loops. In contrast to the all-plus amplitudes, however, they do not have the structure that $F_1^{(2)}$ exhibits; namely, $F_{1;0}^{(2)}$ has uniform transcendental weight two with all other contributions being rational. For the amplitudes with two negative helicities, instead, the FRs have maximum weight two and four at one and two loops, respectively.

4.7 Implementation and performance

The FRs are implemented in the `NJet3 C++` library [1], which is linked to the `PentagonFunctions++` library [278] for the evaluation of the special functions. The six independent helicity amplitudes (shown in Table 4.2) are permuted analytically onto the global basis of pentagon functions defined in the $12 \rightarrow 345$ scattering region to provide the complete list of 16 “mostly-plus” helicity amplitudes required for the sum. This task is performed using the permuted coefficients from the six fully reconstructed amplitudes as an ansatz into the linear relations, so additional reconstruction time is avoided (see Section 4.5.1). Having identified a global basis of pentagon functions for the complete colour and helicity sum, we formulate the partial amplitudes as

$$F^h = c_i^h M_{ij}^h f_j^h, \quad (4.7.1)$$

where: h is the helicity configuration; f_j^h is a list of integers corresponding to the global list of pentagon function monomials, which is evaluated once per phase-space point; M_{ij}^h are sparse matrices of rational numbers that are specific to each partial amplitude; and c_i^h are the independent rational coefficients for each helicity amplitude, written in terms of independent polynomials in the MTVs x_i . The pentagon function monomials are split into parity-odd and -even components, which allows the remaining 16 “mostly-minus” helicities to be computed by simply flipping the parity




Figure 4.4: Flow chart of our evaluation strategy. r/f indicates the rational coefficients and special functions respectively at the three precision levels.

of the special functions and applying complex conjugation to the coefficients. The colour- and helicity-summed ME is constructed numerically from these ingredients. The sparse matrix multiplication is implemented using the `Eigen3` library [269]. Evaluation with `f128s` and `f256s` is provided via the `QD` library [249]. The code is available through `NJet3` [1], where we provide additional installation instructions and example programs demonstrating its usage.

The `C++` code returns the values of the one- and two-loop hard functions, $\mathcal{H}^{(1)}$ and $\mathcal{H}^{(2)}$, obtained by squaring Eq. (4.2.4), substituting the decomposition in N_c and N_f from Eq. (4.2.6), subtracting the IR and UV poles, and finally summing over colour and helicity,

$$\mathcal{H} = \frac{\alpha^2 \alpha_s^3}{(4\pi)^5} \left(\mathcal{H}^{(1)} + \frac{\alpha_s}{4\pi} \mathcal{H}^{(2)} \right) + \mathcal{O}(\alpha_s^5), \quad (4.7.2)$$

$$\mathcal{H}^{(2)} = N_c \mathcal{H}_1^{(2)} + \frac{1}{N_c} \mathcal{H}_2^{(2)} + N_f \mathcal{H}_3^{(2)}.$$

The sum over colours for each helicity can also be returned if required. We find the evaluation time is dominated by the special functions, particularly when higher precision is required. In order to ensure fast and stable numerical evaluation, we adopt the following evaluation strategy, which is also depicted in Fig. 4.4.

1. The user-provided phase-space point is checked for the precision of the on-shell constraints. Points are adjusted in case the precision is not acceptable for the requested number of digits: `f64` \sim 15 digits, `f128` \sim 31 digits, and `f256` \sim 62 digits [249].
2. The colour- and helicity-summed amplitude is computed using `f64` precision at two points which differ only by overall dimension scaling factor. After

accounting for the overall dimension of the squared amplitude, and neglecting scale dependence arising due to the truncation of the perturbative series in α_s to NLO, the two evaluations only differ due to rounding errors at intermediate stages in the evaluation. This accuracy scaling test has been used extensively at one loop. Note that this is unrelated to the scale variation test discussed in Section 1.8.2. We refer to this precision as `f64/f64` since both coefficients and special functions use `f64` precision.

3. If the estimated number of correct digits from the scaling tests falls below a user-defined threshold, the coefficients only are recomputed using `f128` precision after the original point is corrected to `f128` precision (as in step 1). We refer to this as `f128/f64` precision.
4. The scaling test is performed again and if it fails the special functions are reevaluated in `f128` precision. This is `f128/f128` precision.

These steps can be repeated to obtain up to `f256/f256` precision. In practice these steps are rather expensive and unnecessary for standard phenomenological applications, so they are omitted from our strategy.

While the dimension scaling test has been used successfully at one loop, we need to be more careful in our applications when linking the `PentagonFunctions++` library, which also makes use of the dimension rescaling internally. To validate the reliability of the scaling test as an estimate of the error of the result, we evaluate both with a direct `f128/f128` computation and via a scaling test with an error cutoff of three digits at `f64/f64` for a set of 60k points. To ensure a realistic validation, we use “physical” points with a phase-space sampling density determined by the one-loop process, obtained from `NNLOjet`. We compare the estimated error provided by the `f64/f64` scaling test to the relative difference between the `f64/f64` and `f128/f128` evaluations, with the latter taken as the true error. In the following, percentages are always stated with respect to the entire set of points.

As depicted in Fig. 4.5, the scaling test returns a negative for 2.8% of the points. According to true error, an additional 0.2% of the points should be failed and are missed by the scaling test (false positive). Of these points, almost all have true error of three digits, the remaining 0.008% with three digits, so the effect on stability is small. The scaling test also fails some points unnecessarily (false negative), this subset comprising 0.7% of all points, which incurs a small performance penalty in the evaluation strategy. The effects of the false estimates are considered to be allowably small.

We note that the dimension scaling test is statistical and therefore one will always find anomalies in a sufficiently large sample. Care should be taken when integrating




Figure 4.5: Venn diagram of evaluations failing an accuracy test with a cutoff of three digits as estimated by the `f64/f64` scaling test for the right circle (red), and the relative difference between the `f64/f64` and `f128/f128` evaluations for the left circle (blue). The reevaluation error is taken as truth in order to validate the scaling test as a reliable estimate of error. The sets are labelled by their percentage of the entire set of evaluated points in the validation. The leftmost set shows false positives, the centre set shows true negatives, and the rightmost set shows false negatives.

over extreme regions of phase space.

To assess the stability of our implementation (Fig. 4.6) and measure timings, we evaluate the amplitude squared over 100 k points of the physical phase space. We see 1.8% of points failing `f64/f64` evaluation, with 1.2% passing at `f128/f64` and 0.6% passing at `f128/f128`. The evaluation strategy achieves target accuracy for all of the 100 k physical phase-space points tested. We find a single `f64/f64` call has a mean time of 9 s, with 99% of that time spent evaluating the pentagon functions. Using the full evaluation strategy with a target minimum accuracy of three digits, we obtain a mean timing per phase-space point of 26 s.

We present a benchmark evaluation at a point taken from the physical phase space. We choose a generic configuration where the momentum invariants s_{ij} (GeV^2) and pseudoscalar tr_5 (GeV^4) take the values, quoted to four significant figures,

$$\begin{aligned} s_{12} &= 14\,120, & s_{23} &= -1405, & s_{34} &= 7667, \\ s_{45} &= 5493, & s_{15} &= -4404, & \text{tr}_5 &= -17\,600\,000i. \end{aligned} \tag{4.7.3}$$

High precision `f128/f128` evaluations are given in the ancillary files of Ref. [3]. The values for the FRs and the two-loop hard function, normalised by the LO, are shown in Tables 4.2 and 4.3 respectively. The SLC corrections are 600 times smaller than the LC at the point Eq. (4.7.3), while the closed fermion loop corrections are 100




Figure 4.6: Histogram of the error estimate on the two-loop evaluations as given by the scaling test. We use the evaluation strategy with a target accuracy of three digits, denoted by the vertical black line, and show errors for all precision levels as well as the cumulative error on all passing points. A cumulative bin of height h at d digits indicates h points have an accuracy of at best d digits.

Helicity	$N_c F_1^{(2)}/F^{(1)}$	$\frac{1}{N_c} F_2^{(2)}/F^{(1)}$	$N_f F_3^{(2)}/F^{(1)}$
+++++	$-27.76 - 10.17i$	$-1.673 - 0.2396i$	$-5.228 - 4.034i$
-++++	$-25.76 + 27.83i$	$0.3571 - 0.3213i$	$0.3363 - 4.424i$
+++ - +	$-24.16 + 14.59i$	$0.3698 - 0.5539i$	$-4.951 + 0.6672i$
--+++	$-20.23 + 0.8204i$	$-0.4055 - 0.3549i$	$0.053\ 55 + 0.000\ 247\ 8i$
-+-+ +	$-28.58 + 32.90i$	$0.3917 - 0.000\ 548\ 9i$	$3.022 + 1.475i$
+++ --	$-20.94 - 15.34i$	$-0.3080 - 0.4558i$	$-4.880 - 0.005\ 862i$

Table 4.2: Numerical values of the partial amplitudes for the six independent helicities at the benchmark point in Eq. (4.7.3). Values are quoted with $N_c = 3$ and $N_f = 5$, to four significant figures.

$N_c \mathcal{H}_1^{(2)}/\mathcal{H}^{(1)}$	$\frac{1}{N_c} \mathcal{H}_2^{(2)}/\mathcal{H}^{(1)}$	$N_f \mathcal{H}_3^{(2)}/\mathcal{H}^{(1)}$
52.75	0.081 76	0.3956

Table 4.3: Numerical values for the components of the two-loop hard function normalised to the one-loop hard function defined in Eq. (4.7.2) at the benchmark point of Eq. (4.7.3). Values are quoted with $N_c = 3$ and $N_f = 5$, to four significant figures.

times smaller. These ratios do change as we sample different points. Averaging over 100 physical points, the ratio is

$$\left| N_c \frac{\mathcal{H}_1^{(2)}}{\mathcal{H}^{(1)}} \right| : \left| \frac{1}{N_c} \frac{\mathcal{H}_2^{(2)}}{\mathcal{H}^{(1)}} \right| : \left| N_f \frac{\mathcal{H}_3^{(2)}}{\mathcal{H}^{(1)}} \right| = 2000 : 1 : 10.$$

While the evaluation is considerably more difficult than the massless planar five-gluon scattering (Chapter 6), owing to the more complicated set of pentagon functions arising from the non-planar integral topologies, our tests show the amplitudes are clearly ready for phenomenological applications (Chapter 5). In addition, recent improvements to `PentagonFunctions++` [354] would improve the timing and stability, as found in Chapter 6.

4.8 Summary

In this chapter, we have presented a complete FC five-point amplitude at two loops in QCD. All helicity configurations have been implemented in the `NJet3 C++` library, which provides efficient and stable evaluation over the physical scattering region. Although the algebraic complexity of the amplitude is considerable, the direct analytic reconstruction of the FRs was possible by making use of linear relations amongst the coefficients, and univariate partial fractioning that could be done without any analytic knowledge of the intermediate steps in the reduction. We expect these techniques will have applications to other important high-multiplicity

two-loop calculations with more external scales, such as five-particle scattering with an off-shell leg, for which there has also been recent progress [280, 327, 354–365]. We have found a form that is suitable for phenomenological applications. New high-precision predictions for diphoton-plus-jet production at the LHC, which include the dominant N^3LO corrections we have computed here, are presented in Chapter 5.

Chapter 5

NLO QCD corrections to $gg \rightarrow \gamma\gamma g$ distributions

Having obtained the virtual QCD corrections to the amplitude in Chapter 4, we now compute the NLO QCD corrections to the cross section of the gluon-fusion contribution to diphoton-plus-jet production at the LHC. We produce fully differential distributions by combining the two-loop virtual corrections with the one-loop real corrections using antenna subtraction to cancel IR divergences. We observe significant corrections at NLO which demonstrate the importance of combining these corrections with the quark-induced diphoton-plus-jet channel at NNLO.

This chapter is organised as follows. We first discuss recent developments in the field that motivated this study in Section 5.1. Next, in Section 5.2, we review the computational setup, discussing the amplitude-level ingredients and antenna subtraction method used to cancel IR divergences. We then present results for the NLO corrections to differential cross sections at the 13 TeV LHC in Section 5.3. We study the perturbative convergence in both transverse momentum and mass variables as well angular distributions in rapidity and the Collins-Soper angle. We comment on the computational cost of the calculation in Section 5.4, before drawing our conclusions in Section 5.5.

5.1 Background

As discussed in Section 1.8.3, diphoton-plus-jet production is of high phenomenological relevance. The recently computed NNLO corrections of diphoton-plus-jet production [211] display a good perturbative convergence, except in regions where the loop-mediated gluon-fusion subprocess (which contributes to the cross section only from NNLO onwards) is numerically sizeable compared to other contributions. In order to capture the full effects of the QCD corrections, it is important to include

loop-induced gluon-fusion channels from at least one order higher in the perturbative series. These corrections, which are N³LO in the full proton-initiated process, but NLO in the gluon-fusion channel (Section 1.8.1), are the subject of this chapter.

As discussed in Section 2.3, combining and integrating the amplitudes into differential cross sections requires the subtraction of IR divergences. To achieve this in a stable and efficient way is an extremely difficult problem and many solutions have been proposed and applied in calculations up to NNLO. Such subtraction schemes often scale poorly with the number of external particles and only a handful of examples for high-multiplicity processes at NNLO currently exist [211, 319–321].

For the process considered in this chapter, the IR divergences are only at NLO. However, since the real radiation involves two-to-four one-loop squared amplitudes, the automated numerical algorithms for the one-loop amplitudes are tested in extreme phase-space regions. The LO QCD contributions to the gluonic subprocess were first considered in Ref. [252], based on the compact one-loop five-gluon amplitudes [366].

5.2 Computational setup

We consider the inclusive scattering process,

$$gg \rightarrow \gamma\gamma g + X, \quad (5.2.1)$$

at a hadron collider. As the process is loop induced, the LO contribution is at $\mathcal{O}(\alpha_s^3)$ and involves the integration of a one-loop amplitude squared. The NLO QCD corrections are computed by combining the two-loop virtual corrections to the $2 \rightarrow 3$ process $gg \rightarrow \gamma\gamma g$ with the $2 \rightarrow 4$ processes with an additional unresolved parton: $gg \rightarrow \gamma\gamma gg$ and $gg \rightarrow \gamma\gamma q\bar{q}$. Pictorially, we can represent the parton level cross sections up to NLO in QCD as

$$\begin{aligned} \sigma_{gg \rightarrow \gamma\gamma g + X}^{\text{NLO}} = & \int d\Phi_3 \left| \text{Diagram 1} \right|^2 + 2 \int d\Phi_3 \text{Re} \left(\text{Diagram 1}^\dagger \cdot \text{Diagram 2} \right) \\ & + \int d\Phi_4 \left| \text{Diagram 3} \right|^2 + \int d\Phi_4 \left| \text{Diagram 4} \right|^2 + \mathcal{O}(\alpha_s^5), \end{aligned} \quad (5.2.2)$$

where $d\Phi_N$ represents the on-shell phase-space measure for N massless final state particles, Eq. (1.2.5). The one-loop amplitude for $gg \rightarrow \gamma\gamma q\bar{q}$ indicates the loop contribution in which the photons couple to an internal fermion loop. The observable process $pp \rightarrow \gamma\gamma j$ also includes channels where the photons couple to an external

quark pair, which are not included in the gluon-fusion subprocess. The expansion up to the NNLO of $pp \rightarrow \gamma\gamma j$ includes terms up to $\mathcal{O}(\alpha_s^3)$ and so the contributions coming from Eq. (5.2.2) are technically N³LO, as depicted in Fig. 1.6. However, due to the large gluon flux at high-energy hadron colliders, such contributions can be significant (Section 1.8.3).

The one-loop amplitudes for the LO process and the real correction are finite, since the corresponding tree-level processes vanish. The renormalised two-loop five-particle amplitude contains explicit IR divergences generated by the integration over the loop-momenta, while the one-loop six-particle amplitudes exhibit a divergent behaviour when a final-state parton becomes unresolved. The divergences cancel in the final result, as established by the KLN theorem, and an FR of the virtual amplitudes can be defined using QCD factorisation [322]. In our calculation, this cancellation is performed using the antenna subtraction method (Section 2.3). The method extracts the IR singular contributions from the real radiation subprocess, and combines its integrated form with the virtual subprocess, thus enabling their numerical integration using MC methods, performed here in the NNLOjet framework. The QCD structure of the process under consideration is very similar to Higgs-plus-jet production in gluon fusion, which has been computed previously [367, 368] using antenna subtraction, and identical antenna subtraction terms are applied here.

The FRs of the two-loop amplitudes were computed in Chapter 4, providing the FC colour- and helicity-summed expressions through the NJet3 amplitude library. We set a three-digit accuracy threshold for the dimension scaling test (Section 4.7), which guarantees a stable result without significantly affecting the performance.

The one-loop six-particle amplitudes are obtained using a combination of implementations from the OpenLoops2 [369, 370] library and from the generalised unitarity approach [148, 149, 151] within NJet3 (Section 1.6). We use an improved version of OpenLoops2 in combination with the new extension Otter [371], which is a tensor integral library based on the “on-the-fly reduction” [369] of OpenLoops2 and on various stability improvements [370]. This new version of OpenLoops2 allows for a stable computation of the needed one-loop squared amplitudes in deep IR regions. Internally, Otter uses f64 scalar integrals that are provided by Collier [372, 373], as well as f128 scalar integrals provided by OneLoop [255]. Minor modifications were made in NJet3 to avoid de-symmetrisation over the two photons and allow for a pointwise correspondence with the subtraction terms. To compute the one-loop amplitude $gg \rightarrow \gamma\gamma gg$, the OpenLoops2 implementation was generally more efficient, but for exceptional phase-space points it was necessary to use the f128 implementation within NJet3. For the $gg \rightarrow \gamma\gamma\bar{q}q$ channel, we used NJet3, which allowed for a straightforward selection of the required loop contribution. We note that this amplitude is also available within OpenLoops2 and we checked that the

two implementations agree.

The amplitude-level ingredients have been validated in all relevant collinear and soft limits by checking their convergence towards the respective antenna subtraction terms.

5.3 Results

For the numerical evaluation of our NLO results on the gluon-induced diphoton-plus-jet process, we apply the same kinematic cuts as were used for the NNLO calculation of the quark-induced processes [211]. These represent a realistic setup relevant for physics studies at the 13 TeV LHC. Using the quantities defined in Section 2.1.1 and with γ_i as the i^{th} hardest photon, the cuts are as follows.

- Minimum photon transverse momenta:
 - $p_T(\gamma_1) > 30 \text{ GeV}$,
 - $p_T(\gamma_2) > 18 \text{ GeV}$.
- Smooth photon isolation criterion with:
 - $\Delta R = 0.4$,
 - $E_{T,\gamma} = 10 \text{ GeV}$,
 - $\epsilon_\gamma = 1$.
- Minimum pseudorapidity of the photon pair: $|\eta(\gamma\gamma)| < 2.4$.
- Minimum invariant mass of the photon pair: $m(\gamma\gamma) \geq 90 \text{ GeV}$.
- Minimum separation of the photons: $\Delta R(\gamma\gamma) > 0.4$.
- Minimum transverse momentum of the photon pair: $p_T(\gamma\gamma) > 20 \text{ GeV}$.

We consider kinematic distributions in the following diphoton variables.

- Transverse momentum of the diphoton system $p_T(\gamma\gamma)$.
- Diphoton invariant mass $m(\gamma\gamma)$.
- Diphoton total rapidity $|y(\gamma\gamma)|$.
- Diphoton rapidity difference $\Delta y(\gamma\gamma)$.

- Collins-Soper angle¹ $|\phi_{CS}(\gamma\gamma)|$.
- Diphoton azimuthal decorrelation $\Delta\phi(\gamma\gamma)$.

For these distributions, no jet requirement is applied since the transverse momentum cut on the diphoton system—which the jet(s) recoil against—is already sufficient to avoid NNLO-like configurations where all final-state QCD partons become unresolved. This treatment follows that of Ref. [211].

Our numerical results use the NNLO set of the NNPDF3.1 PDFs [273] throughout, thus allowing a straightforward comparison with the existing NNLO results [211] in the quark-initiated channels. The strong coupling is evaluated using LHAPDF [272], with $\alpha_s(m_Z) = 0.118$. The EM coupling is set to $\alpha = 1/137.035999139$. The choices of couplings are also made to match Ref. [211]. MC integration errors are below one percent on average and not displayed in the plots.

The uncertainty on our theory predictions is estimated by a seven-point variation of the renormalisation and factorisation scales around a central value (Section 1.8.2), chosen in a dynamical manner on an event-by-event basis to be

$$\mu_F = \mu_R = \frac{m_T}{2} := \frac{1}{2} \sqrt{m^2(\gamma\gamma) + p_T^2(\gamma\gamma)}, \quad (5.3.2)$$

which is typical for diphoton studies [211, 376, 377].

Figures 5.1 to 5.6 display the theory predictions for the different single-differential distributions in the diphoton variables. We observe the NLO corrections to be sizeable, often being comparable in size to the LO predictions. We define the K -factor as the ratio NLO/LO. The corrections are largest at low $p_T(\gamma\gamma)$ (Fig. 5.1) or $m(\gamma\gamma)$ (Fig. 5.2) where $K \sim 2$ and NLO and LO uncertainties fail to overlap. The ratio smoothly decreases towards $K \sim 1.5$ for large $p_T(\gamma\gamma)$ or $m(\gamma\gamma)$, with overlapping scale uncertainty bands above $p_T(\gamma\gamma) = 200$ GeV or $m(\gamma\gamma) = 175$ GeV.

The integrated cross section is dominated by the region of low $p_T(\gamma\gamma)$ or low $m(\gamma\gamma)$, such that distributions that are differential only in geometrical photon variables (Figs. 5.3 to 5.6) display typically near-uniform $K \sim 2$, and no overlap of the LO and NLO scale uncertainty bands. Visually, the scale uncertainty bands at NLO and LO appear to be of comparable width in all distributions. However, owing to the large size of the NLO corrections, the relative scale uncertainty is reduced from about 50% at LO to 30% at NLO.

¹The Collins-Soper angle [374, 375] $\phi_{CS}(ab)$ for particles a and b is defined by

$$\cos(\phi_{CS}(ab)) = \text{sign}(p_z(ab)) 2 \frac{p_a^+ p_b^- - p_a^- p_b^+}{m(ab) \sqrt{m^2(ab) + p_T^2(ab)}}, \quad p_i^\pm = \frac{1}{\sqrt{2}} (E_i \pm (p_i)_z). \quad (5.3.1)$$




Figure 5.1: Differential distributions in the transverse momentum $p_T(\gamma\gamma)$ of the diphoton system.




Figure 5.2: Differential distributions in the invariant mass $m(\gamma\gamma)$ of the diphoton system.




Figure 5.3: Differential distribution in the Collins-Soper angle $|\cos \phi_{CS}(\gamma\gamma)|$ of the diphoton system.




Figure 5.4: Differential distribution in the azimuthal decorrelation $\Delta\phi(\gamma\gamma)$ of the diphoton system.




Figure 5.5: Differential distribution in the diphoton rapidity difference $\Delta y(\gamma\gamma)$.




Figure 5.6: Differential distribution in the diphoton total rapidity $|y(\gamma\gamma)|$.




Figure 5.7: Two-dimensional differential distribution in the diphoton invariant mass $m(\gamma\gamma)$ and Collins-Soper angle $|\phi_{CS}(\gamma\gamma)|$.




Figure 5.8: Two-dimensional differential distribution in the diphoton rapidity $|y(\gamma\gamma)|$ and transverse momentum $p_T(\gamma\gamma)$.




Figure 5.9: Two-dimensional differential distribution in the diphoton transverse momentum $p_T(\gamma\gamma)$ and invariant mass $m(\gamma\gamma)$, in bins in $p_T(\gamma\gamma)$.




Figure 5.10: Two-dimensional differential distribution in the diphoton transverse momentum $p_T(\gamma\gamma)$ and invariant mass $m(\gamma\gamma)$, for varying lower $p_T(\gamma\gamma)$ cut.

By inspecting the two-dimensional differential distribution in $m(\gamma\gamma)$ and $|\phi_{CS}(\gamma\gamma)|$ (Fig. 5.7) we observe that the relative magnitude of the NLO corrections decreases with increasing $m(\gamma\gamma)$, while the corrections remain uniform in $|\phi_{CS}(\gamma\gamma)|$ for all bins in $m(\gamma\gamma)$. The two-dimensional differential distribution in $|y(\gamma\gamma)|$ and $p_T(\gamma\gamma)$ (Fig. 5.8) also shows the decrease of the corrections towards larger $p_T(\gamma\gamma)$. The decrease is more pronounced at forward rapidity (large $|y(\gamma\gamma)|$) than at central rapidity (small $|y(\gamma\gamma)|$).

The two-dimensional distributions in $p_T(\gamma\gamma)$ and $m(\gamma\gamma)$ (Figs. 5.9 and 5.10) largely reproduce the features of the one-dimensional distributions of Figs. 5.1 and 5.2, both for distributions in bins of $p_T(\gamma\gamma)$ (Fig. 5.9) or for varying lower cut in $p_T(\gamma\gamma)$ (Fig. 5.10). The only novel feature is a non-uniform shape in $m(\gamma\gamma)$ for the highest bin in $p_T(\gamma\gamma)$ (lowest curve in Fig. 5.9), which is indicative of the onset of large logarithmic corrections in $\log(m(\gamma\gamma)/p_T(\gamma\gamma))$ in this range.

The numerical size of the NLO corrections and the scale uncertainties at LO and NLO are comparable to what was observed in inclusive Higgs boson production in gluon fusion [378] or in the Higgs boson transverse momentum distribution in gluon fusion [379, 380]. These processes are mediated through a heavy top quark loop and are very similar to the diphoton-plus-jet production considered here in terms of kinematics and initial-state parton momentum range. The pathology of the NLO corrections observed here is thus not that surprising after all; it does, however, indicate the potential numerical importance of corrections beyond NLO.

The Born-level $gg \rightarrow \gamma\gamma g$ subprocess (corresponding to the LO in our results) contributes to the full diphoton-plus-jet production as part of the NNLO corrections. Corrections to this order were recently computed [211]. They were observed to be moderate and within the scale uncertainty of the previously known NLO results [376] for most of the kinematic range, where they also led to a substantial reduction of the scale uncertainty at NNLO. At low $p_T(\gamma\gamma)$ or low $m(\gamma\gamma)$, larger positive corrections and an increased scale uncertainty were observed [211]. These effects could be identified to be entirely due to the contribution of the $gg \rightarrow \gamma\gamma g$, which only starts to contribute from NNLO onwards, and it was anticipated in Ref. [211] that NLO corrections to the $gg \rightarrow \gamma\gamma g$ (which form a subset of the N³LO corrections to the full diphoton-plus-jet process) could help to stabilise the predictions in the relevant kinematic ranges.

Our results demonstrate that this is not the case. The absolute scale uncertainty on the gluon-induced process does not decrease from LO to NLO, and the NLO correction is of about the size of the LO contribution. Consequently, inclusion of the NLO corrections to the $gg \rightarrow \gamma\gamma g$ into the full NNLO diphoton-plus-jet process will further enhance the predictions at low $p_T(\gamma\gamma)$ or low $m(\gamma\gamma)$, thereby further elongating them from the previously known order, and will leave the scale uncertainty

band largely unchanged. This is an indication that the seven-point scale variation technique may not be providing a reliable estimate of the uncertainty for NNLO diphoton-plus-jet production.

5.4 Computational cost

State-of-the-art calculations in precision phenomenology require high-performance computing resources [381]. In this section, we comment on the computational cost of our calculation.

To reach an average MC error of at most one percent in the total cross section and all differential distributions, we find that approximately 5 M points in the virtual corrections and 30 M points in the real corrections are required after cuts are applied. The real corrections require more evaluations due to the higher dimensionality of the phase space and the presence of IR limits, so a factor of six is unsurprising.

To estimate the relative time per ME call, we time the evaluation of approximately 1 k points after cuts. We find a mean time per point in the virtual corrections of 16 s, with some of the points in `f128` precision taking up to 5 minutes. This is a much shorter mean time than the 26 s found in Section 4.7, which uses the same target accuracy of three digits. Most of this time is in the evaluation of the pentagon functions, the timing of which at fixed precision is sensitive to the degree of IR divergence of the phase-space point. The proportion of points requiring `f128` evaluation, which requires significantly more time than `f64` evaluation, is the dominant contribution to the mean time. Therefore, the timing is highly dependent on the choice of phase space. Both of these benchmarks use version 1.0 of `PentagonFunctions++` [278]; using the recently released version 2.0 [354] would reduce these times, as found in Chapter 6.

For the real contributions, we find an mean time of 11 s for the $gg \rightarrow \gamma\gamma gg$ channel and 14 s for the $gg \rightarrow \gamma\gamma q\bar{q}$ channel. The combined average time per point in the real contributions is 25 s. We stress that the real integration includes extreme regions of phase space, requiring evaluations in `f128` precision.

We find a cost of around 50 k CPU hours in the virtual corrections and 500 k CPU hours in the real corrections for our simulation. The order of magnitude increase in the time cost of the real corrections is good motivation for the NN approximation techniques for high-multiplicity scattering explored in Chapter 3.

5.5 Summary

In this chapter, we have presented the NLO QCD corrections to the diphoton-plus-jet production in the gluon-fusion channel for the first time. The loop-induced process requires the evaluation of six-point one-loop real emission amplitudes and FC five-point two-loop virtual amplitudes. To the best of our knowledge it is the first time that five-point two-loop FC amplitudes have been integrated to provide fully differential cross section predictions relevant for the LHC experiments.

Using a realistic set of kinematic cuts and simulation parameters, we find significant corrections at NLO. This is particularly relevant at low values of $p_T(\gamma\gamma)$ and $m(\gamma\gamma)$. Since observables that are differential only in angles, such as rapidity and the Collins-Soper angle, are inclusive over the energy variables, one observes significant NLO corrections across the full parameter range. Double-differential distributions further highlight this feature, which is reminiscent of the perturbative convergence observed in other gluon-induced processes such as inclusive Higgs production and the Higgs boson transverse momentum distribution. The relative scale uncertainty is reduced by the higher-order corrections, although in absolute terms the scale uncertainty does not decrease from LO to NLO in the low $p_T(\gamma\gamma)$ and $m(\gamma\gamma)$ regions.

This work demonstrates the importance of a combined prediction for quark-induced and gluon-induced diphoton-plus-jet signatures for future precision studies at the LHC.

Chapter 6

Leading-colour double-virtual QCD corrections to $pp \rightarrow 3j$

We present an analytic computation of the helicity amplitudes for trijet production at hadron colliders up to two loops in QCD, providing the virtual and VV contributions for two-to-three NNLO predictions. We provide a fast and stable implementation of the colour- and helicity-summed FRs in C++ as part of the version 3.1.0 release of the `NJet3` library [1].

The conventions, including the decompositions in channels, helicities, and the ratio N_f/N_c , and the IR subtraction to FR, follow those of Ref. [316]. The methods of construction and implementation are the same as those described in Chapter 4. We cross-checked our implementation against the benchmark point in Ref. [316], finding exact agreement.

We express the amplitudes with the scale dependence factorised out. This allows for efficient computation of the amplitudes at a single phase-space point with many different scales, as is used for scale variation uncertainty estimates (Section 1.8.2), with only a single evaluation of the scale-independent part of the amplitudes.

We will now proceed with an analysis of the implementation, looking at the numerical stability and evaluation timing over a sample phase space in Section 6.1, and the performance in IR limits in Section 6.2.

6.1 Stability and timing

We perform similar stability tests on the trijet production channels to those in Section 4.7. We generate 100 k phase-space points isotropically with the algorithm from Ref. [254], which is provided in the `NJet3` library. We use the cuts $p_T > 3 \text{ MeV}$, $\eta < 2.8$, and $\Delta R > 0.4$ for all jets, which is likely to be at least as inclusive over phase space as any phenomenological application. We use arbitrary scales for testing,




Figure 6.1: Stability plots for a single channel in each of the $5g$, $2q3g$, $2q2Q1g$, and $4q1g$ two-loop trijet subprocesses. Other channels within each of the subprocesses perform similarly. Run of 100 k points per channel over a randomly sampled uniform five-point phase space. The evaluation strategy is used with a target accuracy of three digits, denoted on the plots by a black vertical line.

Channel	f64/f64		Precision rescue system	
	Time (s)	Pentagons (%)	Time (s)	Pentagons (%)
$gg \rightarrow ggg$	1.39	69	1.89	77
$gg \rightarrow \bar{q}qg$	1.35	91	1.37	91
$qq \rightarrow qgg$	1.34	92	1.57	93
$q\bar{q} \rightarrow ggg$	1.34	93	1.38	93
$\bar{q}Q \rightarrow Q\bar{q}g$	1.14	99	1.16	99
$\bar{q}\bar{Q} \rightarrow \bar{q}\bar{Q}g$	1.36	99	1.39	99
$\bar{q}g \rightarrow \bar{q}Q\bar{Q}$	1.36	99	1.39	99
$\bar{q}q \rightarrow Q\bar{Q}g$	1.14	99	1.14	99
$\bar{q}g \rightarrow \bar{q}q\bar{q}$	1.84	99	1.90	99
$\bar{q}\bar{q} \rightarrow \bar{q}\bar{q}g$	1.82	99	1.94	99
$\bar{q}q \rightarrow q\bar{q}g$	1.71	99	1.77	99

Table 6.1: Mean timing per phase-space point for combined evaluation of Born, virtual, and VV contributions in f64/f64 and evaluation strategy with target accuracy of three digits.

choosing centre-of-mass energy $\sqrt{s_{12}} = 1$ GeV and renormalisation scale $\mu_R = m_Z$ with Z -boson mass $m_Z = 91.188$ GeV. We evaluate all channels over this phase space using the evaluation strategy with a target accuracy of three digits, generating the stability plots shown in Fig. 6.1 and mean timings of Table 6.1.

The stability and evaluation times improve for the simpler channels, which are those with more quarks and less gluons. In Fig. 6.1, the channels with four quarks even show an `f64` bin with around 1% of the points at the maximum possible precision of 16 digits. The $4q1g$ channels are slower as they are each constructed from a pair of independent $2q2Q1g$ channels, although they show similar stability to the $2q2Q1g$ channels.

Comparing the trijet stability plots to that for $gg \rightarrow \gamma\gamma g$, Fig. 4.6, which includes non-planar integrals, fewer points are required to be reevaluated with higher precision, as expected. The evaluation time is also significantly less for the trijet channels than for $gg \rightarrow \gamma\gamma g$, owing to the LC trijet computations containing only planar diagrams.

Compared to Ref. [316], we find a reduced evaluation time. This is most pronounced in the most difficult channel, pure-gluon scattering $gg \rightarrow ggg$, where the time is halved. This is because the former was evaluated with version 1.0 of `PentagonFunctions++` [278], while our setup uses version 2.0, which additionally supports one massive leg [354]. While we should be careful when comparing to the stability plots of Ref. [354] as they use a different phase space, it also appears that the tail of failing points is shorter in our results, suggesting `PentagonFunctions++` version 2.0 also improves stability.

6.2 Infrared performance

We prepare another phase space using the method described in Section 2.4.1, generating a 100-point slice that approaches an IR limit. The four-point seed phase space is generated randomly using the generator discussed in Section 6.1.

We perform this for ten different four-point seeds and plot the mean in Fig. 6.2 to avoid any irregularities that may arise when approaching the limit in an exceptional direction. The lower precision evaluations diverge when they get too close to the collinear limit; for example in the five-gluon channel, this occurs at around $s_{34}/s_{12} = 10^{-7}$. The origin of this numerical divergence lies in the evaluation of the pentagon functions as it is not fixed by the `f128/f64` evaluation. Evaluation in `f128/f128` remains unproblematic deep into the limit. This demonstrates that these amplitudes are suitable not only for integrating over two-to-three VV phase spaces at NNLO, but also the more difficult two-to-two RVV phase space at N³LO.




Figure 6.2: The two-loop ME for a selection of channels over a slice of phase space which drives into an IR configuration that is collinear in the first two outgoing legs. The error band is given by the dimension scaling test. Where it is not visible, the f_{64}/f_{64} (blue) line coincides with the f_{128}/f_{64} (orange) line, which similarly coincides with the f_{128}/f_{128} (green) line.

Chapter 7

Conclusion

To tackle the challenges that face precision tests of the SM, we must leverage novel techniques in mathematics and computation. The first step for bringing theoretical predictions into the realm of percent precision is the further mapping of the phase space of NNLO processes within fixed-order perturbation theory. In particular, in order for these to be applicable to phenomenology, it is vital that they are implemented in an optimally efficient and stable way.

Gluon-initiated diphoton-plus-jet production offers an attractive testing ground for new technology as it is loop-induced, presenting challenges for conventional event generator techniques. It also exhibits a complicated two-to-three two-loop structure including non-planar integrals, while having simpler poles, being at NLO, and colour than similar partonic processes. Moreover, it is a relevant background for interesting phenomenology such as probing the Higgs coupling.

After building up the technology necessary to perform precision QCD calculations, we discussed the IR behaviour of QCD before presenting a library of IR QCD functions up to at most NNLO, which have a variety of uses in the field.

Next, we reviewed the feasibility of using NNs to optimise the evaluation of the ME for cross section calculations for hadron colliders. We used the $gg \rightarrow \gamma\gamma + n \times g$ amplitudes from `NJet3` to train a NN ensemble, and fed it into the MC event generator `Sherpa` to investigate its behaviour within a full hadronic simulation. At six-point, we found that the total simulation time was sped up by a factor of thirty. This method offers a performant way to run high-multiplicity radiative contributions for event generator simulations where conventional techniques are prohibitively slow.

We then computed the FC two-loop amplitudes for $gg \rightarrow \gamma\gamma g$, with an efficient public implementation released in `NJet3`. We found that their LC contribution contains non-planar integrals, which are those with the highest complexity. Such “industrialised” analytical computations of two-loop five-point QCD processes present huge technical challenges. We first needed a basis of special functions offering fast

and stable evaluation over the physical phase space; for massless and single-mass scattering, the pentagon functions library has recently made this possible. Generating the amplitude via colour-ordered diagrams, we reduced tensor integrals to an MI basis which could be expressed in terms of these special functions. This required constructing and solving a very large system of IBP identities. Finally, we obtained the coefficients in an efficient form using FF reconstruction techniques.

Following the virtual $gg \rightarrow \gamma\gamma g$ amplitude calculation, we used the antenna subtraction scheme to combine the virtual result with the one-loop real corrections to obtain NLO QCD differential cross sections. Considering the full process $pp \rightarrow \gamma\gamma j$, the gluon-fusion subprocess enters at N³LO, but the large gluonic PDF means it provides a dominant correction to NNLO. We computed various observable distributions, finding significant corrections that highlight the importance of including the gluon-initiated channels in these predictions.

We also presented an efficient computation of the LC two-loop amplitudes for hadronic trijet production and tested their public implementation within `NJet3`.

This thesis has focused on the calculation of QCD amplitudes at the precision frontier, in particular, the process of diphoton-plus-jets production through gluon fusion. This presented the challenge of handling the algebraic and analytic complexity of higher-order perturbative expressions. Furthermore, assembly of observables from these contributions required careful regularisation of IR behaviour to cancel poles at each order. All amplitudes were implemented as analytical expressions into the public C++ library `NJet3`; such distribution is vital for the collaborative progress of the whole theoretical community. They provide a vital ingredient for theoretical predictions of cross sections at the LHC in the search for deviations from the SM.

Bibliography

- [1] S. Badger, B. Biedermann, R. Moodie, P. Uwer and V. Yundin, *NJet v3*, <https://bitbucket.org/njet/njet>, 2021.
- [2] R. Moodie, *Optimising hadronic collider simulations using amplitude neural networks*, in *20th International Workshop on Advanced Computing and Analysis Techniques in Physics Research: AI Decoded - Towards Sustainable, Diverse, Performant and Effective Scientific Computing*, 2, 2022, 2202.04506.
- [3] S. Badger, C. Brønnum-Hansen, D. Chicherin, T. Gehrmann, H. B. Hartanto, J. Henn et al., *Virtual QCD corrections to gluon-initiated diphoton plus jet production at hadron colliders*, *JHEP* **11** (2021) 083, [2106.08664].
- [4] S. Badger, T. Gehrmann, M. Marcoli and R. Moodie, *Next-to-leading order QCD corrections to diphoton-plus-jet production through gluon fusion at the LHC*, *Phys. Lett. B* **824** (1, 2022) 136802, [2109.12003].
- [5] S. Badger, C. Brønnum-Hansen, H. B. Hartanto, R. Moodie, T. Peraro and S. Zoia. To appear in SciPost Physics.
- [6] T. Aoyama, T. Kinoshita and M. Nio, *Revised and Improved Value of the QED Tenth-Order Electron Anomalous Magnetic Moment*, *Phys. Rev. D* **97** (2018) 036001, [1712.06060].
- [7] L. Morel, Z. Yao, P. Cladé and S. Guellati-Khélifa, *Determination of the fine-structure constant with an accuracy of 81 parts per trillion*, *Nature* **588** (2020) 61–65.
- [8] PARTICLE DATA GROUP collaboration, P. A. Zyla et al., *Review of Particle Physics*, *PTEP* **2020** (2020) 083C01.
- [9] V. D. Shiltsev, *High energy particle colliders: past 20 years, next 20 years and beyond*, *Phys. Usp.* **55** (2012) 965–976, [1205.3087].
- [10] T. Gehrmann and B. Malaescu, *Precision QCD Physics at the LHC*, 2111.02319.

-
- [11] F. Halzen and A. D. Martin, *Quarks and leptons: an introductory course in modern particle physics*. 1984.
- [12] J. Bagger and J. Wess, *Supersymmetry and supergravity*. 1990.
- [13] M. E. Peskin and D. V. Schroeder, *An Introduction to quantum field theory*. Addison-Wesley, Reading, USA, 1995.
- [14] S. Weinberg, *The Quantum theory of fields. Vol. 1: Foundations*. Cambridge University Press, 6, 2005.
- [15] M. Srednicki, *Quantum field theory*. Cambridge University Press, 1, 2007.
- [16] R. K. Ellis, W. J. Stirling and B. R. Webber, *QCD and collider physics*, vol. 8. Cambridge University Press, 2, 2011, 10.1017/CBO9780511628788.
- [17] J. Collins, *Foundations of perturbative QCD*, vol. 32. Cambridge University Press, 11, 2013.
- [18] J. M. Henn and J. C. Plefka, *Scattering Amplitudes in Gauge Theories, Lect. Notes Phys.* **883** (2014) pp.1–195.
- [19] M. D. Schwartz, *Quantum Field Theory and the Standard Model*. Cambridge University Press, 3, 2014.
- [20] H. Elvang and Y.-t. Huang, *Scattering Amplitudes in Gauge Theory and Gravity*. Cambridge University Press, 4, 2015.
- [21] J. Campbell, J. Huston and F. Krauss, *The black book of quantum chromodynamics: a primer for the LHC era*. Oxford University Press, Oxford, 2018.
- [22] L. J. Dixon, *Calculating scattering amplitudes efficiently*, in *Theoretical Advanced Study Institute in Elementary Particle Physics (TASI 95): QCD and Beyond*, pp. 539–584, 1, 1996, hep-ph/9601359.
- [23] L. J. Dixon, *A brief introduction to modern amplitude methods*, in *Proceedings, 2012 European School of High-Energy Physics (ESHEP 2012): La Pommeraye, Anjou, France, June 06-19, 2012*, pp. 31–67, 2014, 1310.5353, DOI.
- [24] P. Skands, *Introduction to QCD*, in *Theoretical Advanced Study Institute in Elementary Particle Physics: Searching for New Physics at Small and Large Scales*, pp. 341–420, 2013, 1207.2389, DOI.

- [25] M. L. Mangano and S. J. Parke, *Multiparton amplitudes in gauge theories*, *Phys. Rept.* **200** (1991) 301–367, [hep-th/0509223].
- [26] H. K. Dreiner, H. E. Haber and S. P. Martin, *Two-component spinor techniques and Feynman rules for quantum field theory and supersymmetry*, *Phys. Rept.* **494** (2010) 1–196, [0812.1594].
- [27] H. Elvang and Y.-t. Huang, *Scattering Amplitudes*, 1308.1697.
- [28] F. Caola, W. Chen, C. Duhr, X. Liu, B. Mistlberger, F. Petriello et al., *The Path forward to N^3LO* , 3, 2022, 2203.06730.
- [29] G. De Laurentis, *Numerical techniques for analytical high-multiplicity scattering amplitudes*, Ph.D. thesis, Durham U., 2020.
- [30] H. A. Chawdhry, *Higher-Order Calculations in Quantum Chromodynamics*, Ph.D. thesis, University of Cambridge, Cambridge U., 2020. 10.17863/CAM.70435.
- [31] J. C. Whitehead, *The Production of Pairs of Isolated Photons at Higher Orders in QCD*, Ph.D. thesis, Durham U., 2021.
- [32] F. Englert and R. Brout, *Broken Symmetry and the Mass of Gauge Vector Mesons*, *Phys. Rev. Lett.* **13** (1964) 321–323.
- [33] P. W. Higgs, *Broken Symmetries and the Masses of Gauge Bosons*, *Phys. Rev. Lett.* **13** (1964) 508–509.
- [34] J. C. Romao and J. P. Silva, *A resource for signs and Feynman diagrams of the Standard Model*, *Int. J. Mod. Phys. A* **27** (2012) 1230025, [1209.6213].
- [35] M. Gell-Mann, *Symmetries of baryons and mesons*, *Phys. Rev.* **125** (1962) 1067–1084.
- [36] L. D. Faddeev and V. N. Popov, *Feynman Diagrams for the Yang-Mills Field*, *Phys. Lett. B* **25** (1967) 29–30.
- [37] K. G. Wilson and J. B. Kogut, *The Renormalization group and the epsilon expansion*, *Phys. Rept.* **12** (1974) 75–199.
- [38] C. G. Callan, Jr., *Broken scale invariance in scalar field theory*, *Phys. Rev. D* **2** (1970) 1541–1547.
- [39] K. Symanzik, *Small distance behavior in field theory and power counting*, *Commun. Math. Phys.* **18** (1970) 227–246.

- [40] P. A. Baikov, K. G. Chetyrkin and J. H. Kühn, *Five-Loop Running of the QCD coupling constant*, *Phys. Rev. Lett.* **118** (2017) 082002, [1606.08659].
- [41] A. Pich, *Precision physics with inclusive QCD processes*, *Prog. Part. Nucl. Phys.* **117** (2021) 103846, [2012.04716].
- [42] D. J. Gross and F. Wilczek, *Ultraviolet Behavior of Nonabelian Gauge Theories*, *Phys. Rev. Lett.* **30** (1973) 1343–1346.
- [43] H. D. Politzer, *Reliable Perturbative Results for Strong Interactions?*, *Phys. Rev. Lett.* **30** (1973) 1346–1349.
- [44] J. C. Collins, D. E. Soper and G. F. Sterman, *Factorization of Hard Processes in QCD*, *Adv. Ser. Direct. High Energy Phys.* **5** (1989) 1–91, [hep-ph/0409313].
- [45] S. Bailey, T. Cridge, L. A. Harland-Lang, A. D. Martin and R. S. Thorne, *Parton distributions from LHC, HERA, Tevatron and fixed target data: MSHT20 PDFs*, *Eur. Phys. J. C* **81** (2021) 341, [2012.04684].
- [46] J. Gao, L. Harland-Lang and J. Rojo, *The Structure of the Proton in the LHC Precision Era*, *Phys. Rept.* **742** (2018) 1–121, [1709.04922].
- [47] J. J. Ethier and E. R. Nocera, *Parton Distributions in Nucleons and Nuclei*, *Ann. Rev. Nucl. Part. Sci.* **70** (2020) 43–76, [2001.07722].
- [48] S. Amoroso et al., *Snowmass 2021 whitepaper: Proton structure at the precision frontier*, in *2022 Snowmass Summer Study*, 3, 2022, 2203.13923.
- [49] R. D. Ball et al., *The Path to Proton Structure at One-Percent Accuracy*, 2109.02653.
- [50] T.-J. Hou et al., *New CTEQ global analysis of quantum chromodynamics with high-precision data from the LHC*, *Phys. Rev. D* **103** (2021) 014013, [1912.10053].
- [51] H1, ZEUS collaboration, H. Abramowicz et al., *Combination of measurements of inclusive deep inelastic $e^\pm p$ scattering cross sections and QCD analysis of HERA data*, *Eur. Phys. J. C* **75** (2015) 580, [1506.06042].
- [52] S. Alekhin, J. Blümlein, S. Moch and R. Placakyte, *Parton distribution functions, α_s , and heavy-quark masses for LHC Run II*, *Phys. Rev. D* **96** (2017) 014011, [1701.05838].

- [53] M. Constantinou et al., *Lattice QCD Calculations of Parton Physics*, 2202.07193.
- [54] J. M. Campbell et al., *Event Generators for High-Energy Physics Experiments*, 3, 2022, 2203.11110.
- [55] C. Reuschle, *Event simulation for colliders — A basic overview*, *J. Phys. Conf. Ser.* **608** (2015) 012072, [1411.7321].
- [56] SHERPA collaboration, E. Bothmann et al., *Event Generation with Sherpa 2.2*, *SciPost Phys.* **7** (2019) 034, [1905.09127].
- [57] T. Gleisberg, S. Hoeche, F. Krauss, M. Schonherr, S. Schumann, F. Siegert et al., *Event generation with SHERPA 1.1*, *JHEP* **02** (2009) 007, [0811.4622].
- [58] M. Bahr et al., *Herwig++ Physics and Manual*, *Eur. Phys. J. C* **58** (2008) 639–707, [0803.0883].
- [59] G. Corcella, I. G. Knowles, G. Marchesini, S. Moretti, K. Odagiri, P. Richardson et al., *HERWIG 6: An Event generator for hadron emission reactions with interfering gluons (including supersymmetric processes)*, *JHEP* **01** (2001) 010, [hep-ph/0011363].
- [60] J. Bellm et al., *Herwig 7.0/Herwig++ 3.0 release note*, *Eur. Phys. J. C* **76** (2016) 196, [1512.01178].
- [61] T. Sjostrand, S. Mrenna and P. Z. Skands, *PYTHIA 6.4 Physics and Manual*, *JHEP* **05** (2006) 026, [hep-ph/0603175].
- [62] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten et al., *An introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159–177, [1410.3012].
- [63] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer et al., *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, *JHEP* **07** (2014) 079, [1405.0301].
- [64] M. Beneke, *Renormalons*, *Phys. Rept.* **317** (1999) 1–142, [hep-ph/9807443].
- [65] A. Gehrmann-De Ridder, T. Gehrmann and G. Heinrich, *Four particle phase space integrals in massless QCD*, *Nucl. Phys.* **B682** (2004) 265–288, [hep-ph/0311276].

- [66] CMS Collaboration and T. Mc Cauley, “Multi-jet event recorded by the CMS detector (Run 2, 13 TeV).” <https://cds.cern.ch/record/2114784>, 12, 2015.
- [67] G. Luisoni and S. Marzani, *QCD resummation for hadronic final states*, *J. Phys. G* **42** (2015) 103101, [1505.04084].
- [68] M. van Beekveld et al., *Snowmass 2021 White Paper: Resummation for future colliders*, 3, 2022, 2203.07907.
- [69] C. W. Bauer, S. Fleming, D. Pirjol and I. W. Stewart, *An Effective field theory for collinear and soft gluons: Heavy to light decays*, *Phys. Rev. D* **63** (2001) 114020, [hep-ph/0011336].
- [70] C. W. Bauer, D. Pirjol and I. W. Stewart, *Soft collinear factorization in effective field theory*, *Phys. Rev. D* **65** (2002) 054022, [hep-ph/0109045].
- [71] M. Beneke, A. P. Chapovsky, M. Diehl and T. Feldmann, *Soft collinear effective theory and heavy to light currents beyond leading power*, *Nucl. Phys. B* **643** (2002) 431–476, [hep-ph/0206152].
- [72] M. Beneke and T. Feldmann, *Multipole expanded soft collinear effective theory with nonAbelian gauge symmetry*, *Phys. Lett. B* **553** (2003) 267–276, [hep-ph/0211358].
- [73] T. Becher, A. Broggio and A. Ferroglia, *Introduction to Soft-Collinear Effective Theory*, vol. 896. Springer, 2015, 10.1007/978-3-319-14848-9.
- [74] G. C. Fox and S. Wolfram, *A Model for Parton Showers in QCD*, *Nucl. Phys. B* **168** (1980) 285–295.
- [75] M. Bengtsson and T. Sjostrand, *A Comparative Study of Coherent and Noncoherent Parton Shower Evolution*, *Nucl. Phys. B* **289** (1987) 810–846.
- [76] S. Höche, *Introduction to parton-shower event generators*, in *Theoretical Advanced Study Institute in Elementary Particle Physics: Journeys Through the Precision Frontier: Amplitudes for Colliders*, pp. 235–295, 2015, 1411.4085, DOI.
- [77] T. Sjöstrand, *Status and developments of event generators*, *PoS LHCP2016* (2016) 007, [1608.06425].
- [78] V. Bertone and S. Prestel, *Combining N³LO QCD calculations and parton showers for hadronic collision events*, 2202.01082.

- [79] B. Andersson, *The Lund model*, vol. 7. Cambridge University Press, 7, 2005, 10.1017/CBO9780511524363.
- [80] J.-C. Winter, F. Krauss and G. Soff, *A Modified cluster hadronization model*, *Eur. Phys. J. C* **36** (2004) 381–395, [hep-ph/0311085].
- [81] G. P. Salam, *Towards Jetography*, *Eur. Phys. J. C* **67** (2010) 637–686, [0906.1833].
- [82] L. J. Dixon, *Scattering amplitudes: the most perfect microscopic structures in the universe*, *J. Phys. A* **44** (2011) 454001, [1105.0771].
- [83] L. Tancredi, *Computational Techniques in Particle Theory*, in *European Physical Society Conference on High Energy Physics 2021*, 10, 2021, 2111.00205.
- [84] A. B. Goncharov, *Multiple polylogarithms, cyclotomy and modular complexes*, *Math. Res. Lett.* **5** (1998) 497–516, [1105.2076].
- [85] C. Duhr, *Mathematical aspects of scattering amplitudes*, in *Theoretical Advanced Study Institute in Elementary Particle Physics: Journeys Through the Precision Frontier: Amplitudes for Colliders*, pp. 419–476, 2015, 1411.7538, DOI.
- [86] J. L. Bourjaily et al., *Functions Beyond Multiple Polylogarithms for Precision Collider Physics*, 3, 2022, 2203.07088.
- [87] G. Travaglini et al., *The SAGEX Review on Scattering Amplitudes*, 2203.13011.
- [88] G. 't Hooft and M. J. G. Veltman, *Regularization and Renormalization of Gauge Fields*, *Nucl. Phys. B* **44** (1972) 189–213.
- [89] C. Gneidiger et al., *To d, or not to d: recent developments and comparisons of regularization schemes*, *Eur. Phys. J. C* **77** (2017) 471, [1705.01827].
- [90] M. L. Mangano, S. J. Parke and Z. Xu, *Duality and Multi - Gluon Scattering*, *Nucl. Phys. B* **298** (1988) 653–672.
- [91] Z. Bern and D. A. Kosower, *Color decomposition of one loop amplitudes in gauge theories*, *Nucl. Phys. B* **362** (1991) 389–448.
- [92] Z. Bern, L. J. Dixon and D. A. Kosower, *Progress in one loop QCD computations*, *Ann. Rev. Nucl. Part. Sci.* **46** (1996) 109–148, [hep-ph/9602280].

- [93] V. Del Duca, L. J. Dixon and F. Maltoni, *New color decompositions for gauge amplitudes at tree and loop level*, *Nucl. Phys.* **B571** (2000) 51–70, [hep-ph/9910563].
- [94] V. Del Duca, A. Frizzo and F. Maltoni, *Factorization of tree QCD amplitudes in the high-energy limit and in the collinear limit*, *Nucl. Phys. B* **568** (2000) 211–262, [hep-ph/9909464].
- [95] R. Kleiss and H. Kuijf, *Multi - Gluon Cross-sections and Five Jet Production at Hadron Colliders*, *Nucl. Phys. B* **312** (1989) 616–644.
- [96] A. Ochirov and B. Page, *Multi-Quark Colour Decompositions from Unitarity*, *JHEP* **10** (2019) 058, [1908.02695].
- [97] H. Johansson and A. Ochirov, *Color-Kinematics Duality for QCD Amplitudes*, *JHEP* **01** (2016) 170, [1507.00332].
- [98] T. Melia, *Proof of a new colour decomposition for QCD amplitudes*, *JHEP* **12** (2015) 107, [1509.03297].
- [99] C. Reuschle and S. Weinzierl, *Decomposition of one-loop QCD amplitudes into primitive amplitudes based on shuffle relations*, *Phys. Rev. D* **88** (2013) 105020, [1310.0413].
- [100] M. Sjö Dahl, *ColorFull – a C++ library for calculations in $SU(N_c)$ color space*, *Eur. Phys. J. C* **75** (2015) 236, [1412.3967].
- [101] Z. Bern, J. J. M. Carrasco and H. Johansson, *New Relations for Gauge-Theory Amplitudes*, *Phys. Rev. D* **78** (2008) 085011, [0805.3993].
- [102] R. W. Brown and S. G. Naculich, *Color-factor symmetry and BCJ relations for QCD amplitudes*, *JHEP* **11** (2016) 060, [1608.05291].
- [103] S. Keppeler and M. Sjö Dahl, *Orthogonal multiplet bases in $SU(N_c)$ color space*, *JHEP* **09** (2012) 124, [1207.0609].
- [104] Y.-J. Du, M. Sjö Dahl and J. Thorén, *Recursion in multiplet bases for tree-level MHV gluon amplitudes*, *JHEP* **05** (2015) 119, [1503.00530].
- [105] M. Sjö Dahl and J. Thorén, *QCD multiplet bases with arbitrary parton ordering*, *JHEP* **11** (2018) 198, [1809.05002].
- [106] G. 't Hooft, *A Planar Diagram Theory for Strong Interactions*, *Nucl. Phys. B* **72** (1974) 461.

- [107] T. Gleisberg and S. Hoeche, *Comix, a new matrix element generator*, *JHEP* **12** (2008) 039, [0808.3674].
- [108] F. Caravaglios, M. L. Mangano, M. Moretti and R. Pittau, *A New approach to multijet calculations in hadron collisions*, *Nucl. Phys. B* **539** (1999) 215–232, [hep-ph/9807570].
- [109] M. L. Mangano, M. Moretti and R. Pittau, *Multijet matrix elements and shower evolution in hadronic collisions: $Wb\bar{b} + n$ jets as a case study*, *Nucl. Phys. B* **632** (2002) 343–362, [hep-ph/0108069].
- [110] P. Draggiotis, R. H. P. Kleiss and C. G. Papadopoulos, *On the computation of multigluon amplitudes*, *Phys. Lett. B* **439** (1998) 157–164, [hep-ph/9807207].
- [111] P. D. Draggiotis, R. H. P. Kleiss and C. G. Papadopoulos, *Multijet production in hadron collisions*, *Eur. Phys. J. C* **24** (2002) 447–458, [hep-ph/0202201].
- [112] F. Maltoni, K. Paul, T. Stelzer and S. Willenbrock, *Color Flow Decomposition of QCD Amplitudes*, *Phys. Rev. D* **67** (2003) 014026, [hep-ph/0209271].
- [113] C. Duhr, S. Hoeche and F. Maltoni, *Color-dressed recursive relations for multi-parton amplitudes*, *JHEP* **08** (2006) 062, [hep-ph/0607057].
- [114] A. Lifson, C. Reuschle and M. Sjö Dahl, *Introducing the Chirality-flow Formalism*, *Acta Phys. Polon. B* **51** (2020) 1547–1557.
- [115] A. Lifson, C. Reuschle and M. Sjö Dahl, *The chirality-flow formalism*, *Eur. Phys. J. C* **80** (2020) 1006, [2003.05877].
- [116] J. Alnefjord, A. Lifson, C. Reuschle and M. Sjö Dahl, *A Brief Look at the Chirality-Flow Formalism for Standard Model Amplitudes*, *PoS LHCP2021* (2021) 160, [2110.04125].
- [117] J. Alnefjord, A. Lifson, C. Reuschle and M. Sjö Dahl, *The chirality-flow formalism for the standard model*, *Eur. Phys. J. C* **81** (2021) 371, [2011.10075].
- [118] A. Lifson, M. Sjö Dahl and Z. Wettersten, *Automating scattering amplitudes with chirality flow*, 2203.13618.
- [119] R. Frederix and T. Vitos, *The colour matrix at next-to-leading-colour accuracy for tree-level multi-parton processes*, *JHEP* **12** (2021) 157, [2109.10377].
- [120] R. J. Eden, P. V. Landshoff, D. I. Olive and J. C. Polkinghorne, *The analytic S-matrix*. Cambridge Univ. Press, Cambridge, 1966.

- [121] J. Henn, T. Peraro, Y. Xu and Y. Zhang, *A first look at the function space for planar two-loop six-particle Feynman integrals*, 2112.10605.
- [122] E. P. Wigner, *On Unitary Representations of the Inhomogeneous Lorentz Group*, *Annals Math.* **40** (1939) 149–204.
- [123] S. J. Parke and T. R. Taylor, *An Amplitude for n Gluon Scattering*, *Phys. Rev. Lett.* **56** (1986) 2459.
- [124] F. A. Berends and W. T. Giele, *Recursive Calculations for Processes with n Gluons*, *Nucl. Phys.* **B306** (1988) 759–808.
- [125] S. Badger, B. Biedermann and P. Uwer, *One-Loop Amplitudes for Multi-Jet Production at Hadron Colliders*, *PoS RADCOR2011* (2011) 020, [1201.1187].
- [126] R. Britto, F. Cachazo and B. Feng, *New recursion relations for tree amplitudes of gluons*, *Nucl. Phys. B* **715** (2005) 499–522, [hep-th/0412308].
- [127] R. Britto, F. Cachazo, B. Feng and E. Witten, *Direct proof of tree-level recursion relation in Yang-Mills theory*, *Phys. Rev. Lett.* **94** (2005) 181602, [hep-th/0501052].
- [128] S. Badger, B. Biedermann, L. Hackl, J. Plefka, T. Schuster and P. Uwer, *Comparing efficient computation methods for massless QCD tree amplitudes: Closed analytic formulas versus Berends-Giele recursion*, *Phys. Rev. D* **87** (2013) 034011, [1206.2381].
- [129] G. Laurentis and D. Maître, *Extracting analytical one-loop amplitudes from numerical evaluations*, *JHEP* **07** (2019) 123, [1904.04067].
- [130] A. Hodges, *Eliminating spurious poles from gauge-theoretic amplitudes*, *JHEP* **05** (2013) 135, [0905.1473].
- [131] S. Badger, H. Frellesvig and Y. Zhang, *A Two-Loop Five-Gluon Helicity Amplitude in QCD*, *JHEP* **12** (2013) 045, [1310.1051].
- [132] W. T. Giele, Z. Kunszt and K. Melnikov, *Full one-loop amplitudes from tree amplitudes*, *JHEP* **04** (2008) 049, [0801.2237].
- [133] C. F. Berger, Z. Bern, L. J. Dixon, F. Febres Cordero, D. Forde, H. Ita et al., *An Automated Implementation of On-Shell Methods for One-Loop Amplitudes*, *Phys. Rev. D* **78** (2008) 036003, [0803.4180].

- [134] G. Bevilacqua, M. Czakon, M. V. Garzelli, A. van Hameren, A. Kardos, C. G. Papadopoulos et al., *HELAC-NLO*, *Comput. Phys. Commun.* **184** (2013) 986–997, [1110.1499].
- [135] V. Hirschi, R. Frederix, S. Frixione, M. V. Garzelli, F. Maltoni and R. Pittau, *Automation of one-loop QCD corrections*, *JHEP* **05** (2011) 044, [1103.0621].
- [136] F. Cascioli, P. Maierhofer and S. Pozzorini, *Scattering Amplitudes with Open Loops*, *Phys. Rev. Lett.* **108** (2012) 111601, [1111.5206].
- [137] G. Cullen et al., *GOSAM-2.0: a tool for automated one-loop calculations within the Standard Model and beyond*, *Eur. Phys. J. C* **74** (2014) 3001, [1404.7096].
- [138] A. Denner, J.-N. Lang and S. Uccirati, *Recola2: REcursive Computation of One-Loop Amplitudes 2*, *Comput. Phys. Commun.* **224** (2018) 346–361, [1711.07388].
- [139] C. Degrande, V. Hirschi and O. Mattelaer, *Automated Computation of One-Loop Amplitudes*, *Ann. Rev. Nucl. Part. Sci.* **68** (2018) 291–312.
- [140] S. Pozzorini, N. Schär and M. F. Zoller, *Two-loop tensor integral coefficients in OpenLoops*, 2201.11615.
- [141] S. Pozzorini, N. Schär and M. F. Zoller, *Two-loop amplitude generation in OpenLoops*, in *15th International Symposium on Radiative Corrections: Applications of Quantum Field Theory to Phenomenology AND LoopFest XIX: Workshop on Radiative Corrections for the LHC and Future Colliders*, 2, 2022, 2202.00703.
- [142] R. K. Ellis, Z. Kunszt, K. Melnikov and G. Zanderighi, *One-loop calculations in quantum field theory: from Feynman diagrams to unitarity cuts*, *Phys. Rept.* **518** (2012) 141–250, [1105.4319].
- [143] Z. Bern and Y.-t. Huang, *Basics of Generalized Unitarity*, *J. Phys. A* **44** (2011) 454003, [1103.1869].
- [144] R. Britto, *Loop Amplitudes in Gauge Theories: Modern Analytic Approaches*, *J. Phys. A* **44** (2011) 454006, [1012.4493].
- [145] C. F. Berger and D. Forde, *Multi-Parton Scattering Amplitudes via On-Shell Methods*, *Ann. Rev. Nucl. Part. Sci.* **60** (2010) 181–205, [0912.3534].
- [146] S. Weinzierl, *Feynman Integrals*, 2201.03593.

- [147] R. E. Cutkosky, *Singularities and discontinuities of Feynman amplitudes*, *J. Math. Phys.* **1** (1960) 429–433.
- [148] Z. Bern, L. J. Dixon, D. C. Dunbar and D. A. Kosower, *One loop n point gauge theory amplitudes, unitarity and collinear limits*, *Nucl. Phys. B* **425** (1994) 217–260, [hep-ph/9403226].
- [149] Z. Bern, L. J. Dixon, D. C. Dunbar and D. A. Kosower, *Fusing gauge theory tree amplitudes into loop amplitudes*, *Nucl. Phys. B* **435** (1995) 59–101, [hep-ph/9409265].
- [150] R. K. Ellis and G. Zanderighi, *Scalar one-loop integrals for QCD*, *JHEP* **02** (2008) 002, [0712.1851].
- [151] R. Britto, F. Cachazo and B. Feng, *Generalized unitarity and one-loop amplitudes in $N=4$ super-Yang-Mills*, *Nucl. Phys. B* **725** (2005) 275–305, [hep-th/0412103].
- [152] N. Arkani-Hamed, F. Cachazo and J. Kaplan, *What is the Simplest Quantum Field Theory?*, *JHEP* **09** (2010) 016, [0808.1446].
- [153] R. Britto and B. Feng, *Integral coefficients for one-loop amplitudes*, *JHEP* **02** (2008) 095, [0711.4284].
- [154] D. Forde, *Direct extraction of one-loop integral coefficients*, *Phys. Rev. D* **75** (2007) 125019, [0704.1835].
- [155] C. Anastasiou, R. Britto, B. Feng, Z. Kunszt and P. Mastrolia, *D -dimensional unitarity cut method*, *Phys. Lett. B* **645** (2007) 213–216, [hep-ph/0609191].
- [156] C. Anastasiou, R. Britto, B. Feng, Z. Kunszt and P. Mastrolia, *Unitarity cuts and Reduction to master integrals in d dimensions for one-loop amplitudes*, *JHEP* **03** (2007) 111, [hep-ph/0612277].
- [157] S. D. Badger, *Direct Extraction Of One Loop Rational Terms*, *JHEP* **01** (2009) 049, [0806.4600].
- [158] W. L. van Neerven and J. A. M. Vermaseren, *Large loop integrals*, *Phys. Lett. B* **137** (1984) 241–244.
- [159] G. Ossola, C. G. Papadopoulos and R. Pittau, *Reducing full one-loop amplitudes to scalar integrals at the integrand level*, *Nucl. Phys. B* **763** (2007) 147–169, [hep-ph/0609007].

- [160] K. G. Chetyrkin and F. V. Tkachov, *Integration by Parts: The Algorithm to Calculate beta Functions in 4 Loops*, *Nucl. Phys.* **B192** (1981) 159–204.
- [161] A. G. Grozin, *Integration by parts: An Introduction*, *Int. J. Mod. Phys. A* **26** (2011) 2807–2854, [1104.3993].
- [162] F. V. Tkachov, *A Theorem on Analytical Calculability of Four Loop Renormalization Group Functions*, *Phys. Lett.* **100B** (1981) 65–68.
- [163] S. Laporta, *High precision calculation of multiloop Feynman integrals by difference equations*, *Int. J. Mod. Phys. A* **15** (2000) 5087–5159, [hep-ph/0102033].
- [164] J. Klappert and F. Lange, *Reconstructing rational functions with FireFly*, *Comput. Phys. Commun.* **247** (2020) 106951, [1904.00009].
- [165] T. Peraro, *Scattering amplitudes over finite fields and multivariate functional reconstruction*, *JHEP* **12** (2016) 030, [1608.01902].
- [166] A. von Manteuffel and R. M. Schabinger, *A novel approach to integration by parts reduction*, *Phys. Lett.* **B744** (2015) 101–104, [1406.4513].
- [167] T. Peraro, *FiniteFlow: multivariate functional reconstruction using finite fields and dataflow graphs*, *JHEP* **07** (2019) 031, [1905.08019].
- [168] A. V. Smirnov and F. S. Chuharev, *FIRE6: Feynman Integral REduction with Modular Arithmetic*, *Comput. Phys. Commun.* **247** (2020) 106877, [1901.07808].
- [169] J. Klappert, S. Y. Klein and F. Lange, *Interpolation of dense and sparse rational functions and other improvements in FireFly*, *Comput. Phys. Commun.* **264** (2021) 107968, [2004.01463].
- [170] J. Klappert, F. Lange, P. Maierhöfer and J. Usovitsch, *Integral Reduction with Kira 2.0 and Finite Field Methods*, *Comput. Phys. Commun.* **266** (2021) 108024, [2008.06494].
- [171] S. Abreu, J. Dormans, F. Febres Cordero, H. Ita, M. Kraus, B. Page et al., *Caravel: A C++ framework for the computation of multi-loop amplitudes with numerical unitarity*, *Comput. Phys. Commun.* **267** (2021) 108069, [2009.11957].
- [172] D. E. Knuth, *The art of computer programming. Vol. 2, Seminumerical algorithms*. Addison-Wesley, Reading, USA, 3rd ed., 1997.

- [173] ATLAS collaboration, G. Aad et al., *The ATLAS Experiment at the CERN Large Hadron Collider*, *JINST* **3** (2008) S08003.
- [174] CMS collaboration, S. Chatrchyan et al., *The CMS Experiment at the CERN LHC*, *JINST* **3** (2008) S08004.
- [175] ATLAS collaboration, *Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC*, tech. rep., CERN, Geneva, 6, 2019.
- [176] CMS collaboration, A. M. Sirunyan et al., *Precision luminosity measurement in proton-proton collisions at $\sqrt{s} = 13$ TeV in 2015 and 2016 at CMS*, *Eur. Phys. J. C* **81** (2021) 800, [2104.01927].
- [177] CMS collaboration, V. Khachatryan et al., *Jet energy scale and resolution in the CMS experiment in pp collisions at 8 TeV*, *JINST* **12** (2017) P02014, [1607.03663].
- [178] ATLAS collaboration, M. Aaboud et al., *Jet energy scale measurements and their systematic uncertainties in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev. D* **96** (2017) 072002, [1703.09665].
- [179] G. P. Salam, *Theory vision*, *PoS LHCP2018* (2018) 304, [1811.11282].
- [180] A. Dainese, M. Mangano, A. B. Meyer, A. Nisati, G. Salam and M. A. Vesterinen, eds., *Report on the Physics at the HL-LHC, and Perspectives for the HE-LHC*, vol. 7/2019 of *CERN Yellow Reports: Monographs*. CERN, Geneva, Switzerland, 2019, 10.23731/CYRM-2019-007.
- [181] D. A. Kosower and K. J. Larsen, *Maximal Unitarity at Two Loops*, *Phys. Rev. D* **85** (2012) 045017, [1108.1180].
- [182] P. Mastrolia and G. Ossola, *On the Integrand-Reduction Method for Two-Loop Scattering Amplitudes*, *JHEP* **11** (2011) 014, [1107.6041].
- [183] S. Badger, H. Frellesvig and Y. Zhang, *Hepta-Cuts of Two-Loop Scattering Amplitudes*, *JHEP* **04** (2012) 055, [1202.2019].
- [184] Y. Zhang, *Integrand-Level Reduction of Loop Amplitudes by Computational Algebraic Geometry Methods*, *JHEP* **09** (2012) 042, [1205.5707].
- [185] P. Mastrolia, E. Mirabella, G. Ossola and T. Peraro, *Scattering Amplitudes from Multivariate Polynomial Division*, *Phys. Lett. B* **718** (2012) 173–177, [1205.7087].

- [186] P. Mastrolia, E. Mirabella, G. Ossola and T. Peraro, *Integrand-Reduction for Two-Loop Scattering Amplitudes through Multivariate Polynomial Division*, *Phys. Rev. D* **87** (2013) 085026, [1209.4319].
- [187] H. Ita, *Two-loop Integrand Decomposition into Master Integrals and Surface Terms*, *Phys. Rev. D* **94** (2016) 116015, [1510.05626].
- [188] S. Badger, G. Mogull, A. Ochirov and D. O'Connell, *A Complete Two-Loop, Five-Gluon Helicity Amplitude in Yang-Mills Theory*, *JHEP* **10** (2015) 064, [1507.08797].
- [189] S. Abreu, F. Febres Cordero, H. Ita, M. Jaquier, B. Page and M. Zeng, *Two-Loop Four-Gluon Amplitudes from Numerical Unitarity*, *Phys. Rev. Lett.* **119** (2017) 142001, [1703.05273].
- [190] M. Cacciari and N. Houdeau, *Meaningful characterisation of perturbative theoretical uncertainties*, *JHEP* **09** (2011) 039, [1105.5152].
- [191] M. Bonvini, *Probabilistic definition of the perturbative theoretical uncertainty from missing higher orders*, *Eur. Phys. J. C* **80** (2020) 989, [2006.16293].
- [192] C. Duhr, A. Huss, A. Mazeliauskas and R. Szafron, *An analysis of Bayesian estimates for missing higher orders in perturbative calculations*, *JHEP* **09** (2021) 122, [2106.04585].
- [193] F. James, *Monte Carlo Theory and Practice*, *Rept. Prog. Phys.* **43** (1980) 1145.
- [194] I. Hughes and T. Hase, *Measurements and their Uncertainties: A practical guide to modern error analysis*. OUP Oxford, 2010.
- [195] F. Caola, K. Melnikov and M. Schulze, *Fiducial cross sections for Higgs boson production in association with a jet at next-to-next-to-leading order in QCD*, *Phys. Rev. D* **92** (2015) 074032, [1508.02684].
- [196] R. Franceschini, G. F. Giudice, J. F. Kamenik, M. McCullough, A. Pomarol, R. Rattazzi et al., *What is the $\gamma\gamma$ resonance at 750 GeV?*, *JHEP* **03** (2016) 144, [1512.04933].
- [197] S. Catani, L. Cieri, D. de Florian, G. Ferrera and M. Grazzini, *Diphoton production at hadron colliders: a fully-differential QCD calculation at NNLO*, *Phys. Rev. Lett.* **108** (2012) 072001, [1110.2375].

- [198] J. M. Campbell, R. K. Ellis, Y. Li and C. Williams, *Predictions for diphoton production at the LHC through NNLO in QCD*, *JHEP* **07** (2016) 148, [1603.02663].
- [199] T. Gehrmann, N. Glover, A. Huss and J. Whitehead, *Scale and isolation sensitivity of diphoton distributions at the LHC*, *JHEP* **01** (2021) 108, [2009.11310].
- [200] C. Anastasiou, E. W. N. Glover and M. E. Tejeda-Yeomans, *Two loop QED and QCD corrections to massless fermion boson scattering*, *Nucl. Phys. B* **629** (2002) 255–289, [hep-ph/0201274].
- [201] Z. Bern, A. De Freitas and L. J. Dixon, *Two loop amplitudes for gluon fusion into two photons*, *JHEP* **09** (2001) 037, [hep-ph/0109078].
- [202] T. Neumann, *The diphoton q_T spectrum at $N^3LL' + NNLO$* , *Eur. Phys. J. C* **81** (2021) 905, [2107.12478].
- [203] F. Caola, A. Von Manteuffel and L. Tancredi, *Diphoton Amplitudes in Three-Loop Quantum Chromodynamics*, *Phys. Rev. Lett.* **126** (2021) 112004, [2011.13946].
- [204] P. Bargiela, F. Caola, A. von Manteuffel and L. Tancredi, *Three-loop helicity amplitudes for diphoton production in gluon fusion*, *JHEP* **02** (2022) 153, [2111.13595].
- [205] J. R. Andersen et al., *Les Houches 2015: Physics at TeV Colliders Standard Model Working Group Report*, in *9th Les Houches Workshop on Physics at TeV Colliders*, 5, 2016, 1605.04692.
- [206] *Les Houches 2017: Physics at TeV Colliders Standard Model Working Group Report*, 3, 2018.
- [207] S. Amoroso et al., *Les Houches 2019: Physics at TeV Colliders: Standard Model Working Group Report*, in *11th Les Houches Workshop on Physics at TeV Colliders: PhysTeV Les Houches*, 3, 2020, 2003.01700.
- [208] B. Agarwal, F. Buccioni, A. von Manteuffel and L. Tancredi, *Two-loop leading colour QCD corrections to $q\bar{q} \rightarrow \gamma\gamma g$ and $qg \rightarrow \gamma\gamma q$* , *JHEP* **04** (2021) 201, [2102.01820].
- [209] H. A. Chawdhry, M. Czakon, A. Mitov and R. Poncelet, *Two-loop leading-colour QCD helicity amplitudes for two-photon plus jet production at the LHC*, *JHEP* **07** (2021) 164, [2103.04319].

- [210] B. Agarwal, F. Buccioni, A. von Manteuffel and L. Tancredi, *Two-loop helicity amplitudes for diphoton plus jet production in full color*, *Phys. Rev. Lett.* **127** (12, 2021) 262001, [2105.04585].
- [211] H. A. Chawdhry, M. Czakon, A. Mitov and R. Poncelet, *NNLO QCD corrections to diphoton production with an additional jet at the LHC*, *JHEP* **09** (9, 2021) 093, [2105.06940].
- [212] ATLAS collaboration, G. Aad et al., *Measurement of the production cross section of pairs of isolated photons in pp collisions at 13 TeV with the ATLAS detector*, *JHEP* **11** (2021) 169, [2107.09330].
- [213] N. Agarwal, L. Magnea, C. Signorile-Signorile and A. Tripathi, *The Infrared Structure of Perturbative Gauge Theories*, 2112.07099.
- [214] M. Cacciari, G. P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, *JHEP* **04** (2008) 063, [0802.1189].
- [215] M. Cacciari, G. P. Salam and G. Soyez, *FastJet User Manual*, *Eur. Phys. J. C* **72** (2012) 1896, [1111.6097].
- [216] S. Frixione, *Isolated photons in perturbative QCD*, *Phys. Lett. B* **429** (1998) 369–374, [hep-ph/9801442].
- [217] T. Kinoshita, *Mass singularities of Feynman amplitudes*, *J. Math. Phys.* **3** (1962) 650–677.
- [218] T. D. Lee and M. Nauenberg, *Degenerate Systems and Mass Singularities*, *Phys. Rev.* **133** (1964) B1549–B1562.
- [219] D. A. Kosower, *Antenna factorization of gauge theory amplitudes*, *Phys. Rev. D* **57** (1998) 5410–5416, [hep-ph/9710213].
- [220] T. G. Birthwright, E. W. N. Glover, V. V. Khoze and P. Marquard, *Collinear limits in QCD from MHV rules*, *JHEP* **07** (2005) 068, [hep-ph/0505219].
- [221] S. Catani and M. Grazzini, *Infrared factorization of tree level QCD amplitudes at the next-to-next-to-leading order and beyond*, *Nucl. Phys. B* **570** (2000) 287–325, [hep-ph/9908523].
- [222] J. M. Campbell and E. W. N. Glover, *Double unresolved approximations to multiparton scattering amplitudes*, *Nucl. Phys. B* **527** (1998) 264–288, [hep-ph/9710255].

- [223] S. Badger, F. Buciuni and T. Peraro, *One-loop triple collinear splitting amplitudes in QCD*, *JHEP* **09** (2015) 188, [1507.05070].
- [224] Z. Bern and G. Chalmers, *Factorization in one loop gauge theory*, *Nucl. Phys. B* **447** (1995) 465–518, [hep-ph/9503236].
- [225] P. Banerjee, T. Engel, N. Schalch, A. Signer and Y. Ulrich, *Bhabha scattering at NNLO with next-to-soft stabilisation*, *Phys. Lett. B* **820** (2021) 136547, [2106.07469].
- [226] S. Frixione, Z. Kunszt and A. Signer, *Three jet cross-sections to next-to-leading order*, *Nucl. Phys. B* **467** (1996) 399–442, [hep-ph/9512328].
- [227] R. Frederix, S. Frixione, F. Maltoni and T. Stelzer, *Automation of next-to-leading order computations in QCD: The FKS subtraction*, *JHEP* **10** (2009) 003, [0908.4272].
- [228] S. Catani and M. H. Seymour, *A General algorithm for calculating jet cross-sections in NLO QCD*, *Nucl. Phys. B* **485** (1997) 291–419, [hep-ph/9605323].
- [229] S. Catani, S. Dittmaier, M. H. Seymour and Z. Trocsanyi, *The Dipole formalism for next-to-leading order QCD calculations with massive partons*, *Nucl. Phys.* **B627** (2002) 189–265, [hep-ph/0201036].
- [230] W. J. Torres Bobadilla et al., *May the four be with you: Novel IR-subtraction methods to tackle NNLO calculations*, *Eur. Phys. J. C* **81** (2021) 250, [2012.02567].
- [231] L. Magnea, G. Pelliccioli, C. Signorile-Signorile, P. Torrielli and S. Uccirati, *Analytic integration of soft and collinear radiation in factorised QCD cross sections at NNLO*, *JHEP* **02** (2021) 037, [2010.14493].
- [232] G. Heinrich, *Collider Physics at the Precision Frontier*, *Phys. Rept.* **922** (2021) 1–69, [2009.00516].
- [233] A. Gehrmann-De Ridder, T. Gehrmann and E. W. N. Glover, *Antenna subtraction at NNLO*, *JHEP* **09** (2005) 056, [hep-ph/0505111].
- [234] A. Daleo, T. Gehrmann and D. Maitre, *Antenna subtraction with hadronic initial states*, *JHEP* **04** (2007) 016, [hep-ph/0612257].
- [235] J. Currie, E. W. N. Glover and S. Wells, *Infrared Structure at NNLO Using Antenna Subtraction*, *JHEP* **04** (2013) 066, [1301.4693].

- [236] X. Chen, T. Gehrmann, N. Glover, A. Huss and M. Marcoli, *Automation of antenna subtraction in colour space: gluonic processes*, 2203.13531.
- [237] M. Czakon, *A novel subtraction scheme for double-real radiation at NNLO*, *Phys. Lett. B* **693** (2010) 259–268, [1005.0274].
- [238] R. Boughezal, K. Melnikov and F. Petriello, *A subtraction scheme for NNLO computations*, *Phys. Rev. D* **85** (2012) 034025, [1111.7041].
- [239] M. Czakon and D. Heymes, *Four-dimensional formulation of the sector-improved residue subtraction scheme*, *Nucl. Phys. B* **890** (2014) 152–227, [1408.2500].
- [240] F. Caola, K. Melnikov and R. Röntsch, *Nested soft-collinear subtractions in NNLO QCD computations*, *Eur. Phys. J. C* **77** (2017) 248, [1702.01352].
- [241] V. Del Duca, C. Duhr, A. Kardos, G. Somogyi, Z. Ször, Z. Trócsányi et al., *Jet production in the CoLoRFulNNLO method: event shapes in electron-positron collisions*, *Phys. Rev. D* **94** (2016) 074019, [1606.03453].
- [242] M. Cacciari, F. A. Dreyer, A. Karlberg, G. P. Salam and G. Zanderighi, *Fully Differential Vector-Boson-Fusion Higgs Production at Next-to-Next-to-Leading Order*, *Phys. Rev. Lett.* **115** (2015) 082002, [1506.02660].
- [243] L. Magnea, E. Maina, G. Pelliccioli, C. Signorile-Signorile, P. Torrielli and S. Uccirati, *Local analytic sector subtraction at NNLO*, *JHEP* **12** (2018) 107, [1806.09570].
- [244] L. Magnea, E. Maina, G. Pelliccioli, C. Signorile-Signorile, P. Torrielli and S. Uccirati, *Factorisation and Subtraction beyond NLO*, *JHEP* **12** (2018) 062, [1809.05444].
- [245] S. Catani and M. Grazzini, *An NNLO subtraction formalism in hadron collisions and its application to Higgs boson production at the LHC*, *Phys. Rev. Lett.* **98** (2007) 222002, [hep-ph/0703012].
- [246] J. Gaunt, M. Stahlhofen, F. J. Tackmann and J. R. Walsh, *N-jettiness Subtractions for NNLO QCD Calculations*, *JHEP* **09** (2015) 058, [1505.04794].
- [247] F. Herzog, *Geometric IR subtraction for final state real radiation*, *JHEP* **08** (2018) 006, [1804.07949].
- [248] D. A. Kosower, *All order collinear behavior in gauge theories*, *Nucl. Phys. B* **552** (1999) 319–336, [hep-ph/9901201].

- [249] Y. Hida, X. S. Li and D. H. Bailey, *libqd: quad-double / double-double computation package*, <https://www.davidhbailey.com/dhbsoftware/>, 2010.
- [250] Z. Bern, V. Del Duca, W. B. Kilgore and C. R. Schmidt, *The infrared behavior of one loop QCD amplitudes at next-to-next-to leading order*, *Phys. Rev. D* **60** (1999) 116001, [[hep-ph/9903516](#)].
- [251] D. A. Dicus and S. S. D. Willenbrock, *Photon Pair Production and the Intermediate Mass Higgs Boson*, *Phys. Rev. D* **37** (1988) 1801.
- [252] D. de Florian and Z. Kunszt, *Two photons plus jet at LHC: The NNLO contribution from the gg initiated process*, *Phys. Lett. B* **460** (1999) 184–188, [[hep-ph/9905283](#)].
- [253] S. Catani, D. de Florian and G. Rodrigo, *The Triple collinear limit of one loop QCD amplitudes*, *Phys. Lett. B* **586** (2004) 323–331, [[hep-ph/0312067](#)].
- [254] E. Byckling and K. Kajantie, *Particle Kinematics: (Chapters I-VI, X)*. University of Jyvaskyla, Jyvaskyla, Finland, 1971.
- [255] A. van Hameren, *OneLOop: For the evaluation of one-loop scalar functions*, *Comput. Phys. Commun.* **182** (2011) 2427–2438, [[1007.4716](#)].
- [256] S. Badger, B. Biedermann, P. Uwer and V. Yundin, *Numerical evaluation of virtual corrections to multi-jet production in massless QCD*, *Comput. Phys. Commun.* **184** (2013) 1981–1998, [[1209.0100](#)].
- [257] S. Badger, B. Biedermann, P. Uwer and V. Yundin, *Computation of multi-leg amplitudes with NJet*, *J. Phys. Conf. Ser.* **523** (2014) 012057, [[1312.7140](#)].
- [258] J. Aylett-Bullock and R. Moodie, *n3jet_diphoton v1*, https://gitlab.com/JosephPB/n3jet_diphoton, 2021.
- [259] M. Feickert and B. Nachman, *A Living Review of Machine Learning for Particle Physics*, [2102.02770](#).
- [260] A. Butter et al., *Machine Learning and LHC Event Generation*, **3**, 2022, [2203.07460](#).
- [261] I. Goodfellow, Y. Bengio and A. Courville, *Deep Learning*. MIT Press, 2016.
- [262] S. Badger and J. Bullock, *Using neural networks for efficient evaluation of high multiplicity scattering amplitudes*, *JHEP* **06** (2020) 114, [[2002.07516](#)].

- [263] J. Aylett-Bullock, S. Badger and R. Moodie, *Optimising simulations for diphoton production at hadron colliders using amplitude neural networks*, *JHEP* **2021** (8, 2021) , [2106.09474].
- [264] F. Chollet et al., *Keras*, <https://github.com/fchollet/keras>, 2015.
- [265] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro et al., *TensorFlow*, <https://www.tensorflow.org/>, 2015.
- [266] D. P. Kingma and J. Ba, *Adam: A method for stochastic optimization*, *3rd International Conference for Learning Representations* (2015) , [1412.6980].
- [267] R. Kleiss, W. J. Stirling and S. D. Ellis, *A New Monte Carlo Treatment of Multiparticle Phase Space at High-energies*, *Comput. Phys. Commun.* **40** (1986) 359.
- [268] T. Ohl, *Vegas revisited: Adaptive Monte Carlo integration beyond factorization*, *Comput. Phys. Commun.* **120** (1999) 13–19, [hep-ph/9806432].
- [269] G. Guennebaud, B. Jacob et al., *Eigen v3*, <https://eigen.tuxfamily.org>, 2010.
- [270] T. Binoth et al., *A Proposal for a Standard Interface between Monte Carlo Tools and One-Loop Programs*, *Comput. Phys. Commun.* **181** (2010) 1612–1622, [1001.1307].
- [271] S. Alioli et al., *Update of the Binoth Les Houches Accord for a standard interface between Monte Carlo tools and one-loop programs*, *Comput. Phys. Commun.* **185** (2014) 560–571, [1308.3462].
- [272] A. Buckley, J. Ferrando, S. Lloyd, K. Nordström, B. Page, M. Rüfenacht et al., *LHAPDF6: parton density access in the LHC precision era*, *Eur. Phys. J. C* **75** (2015) 132, [1412.7420].
- [273] NNPDF collaboration, R. D. Ball et al., *Parton distributions from high-precision collider data*, *Eur. Phys. J. C* **77** (2017) 663, [1706.00428].
- [274] S. Badger, A. Guffanti and V. Yundin, *Next-to-leading order QCD corrections to di-photon production in association with up to three jets at the Large Hadron Collider*, *JHEP* **03** (2014) 122, [1312.5927].
- [275] C. Bierlich et al., *Robust Independent Validation of Experiment and Theory: Rivet version 3*, *SciPost Phys.* **8** (2020) 026, [1912.05451].

- [276] ATLAS collaboration, M. Aaboud et al., *Measurements of integrated and differential cross sections for isolated photon pair production in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *Phys. Rev. D* **95** (2017) 112005, [1704.03839].
- [277] D. Maître and H. Truong, *A factorisation-aware Matrix element emulator*, *JHEP* **11** (2021) 066, [2107.06625].
- [278] D. Chicherin and V. Sotnikov, *Pentagon Functions for Scattering of Five Massless Particles*, *JHEP* **12** (2020) 167, [2009.07803].
- [279] D. Chicherin, J. Henn and V. Mitev, *Bootstrapping pentagon functions*, *JHEP* **05** (2018) 164, [1712.09610].
- [280] C. G. Papadopoulos, D. Tommasini and C. Wever, *The Pentabox Master Integrals with the Simplified Differential Equations approach*, *JHEP* **04** (2016) 078, [1511.09404].
- [281] T. Gehrmann, J. Henn and N. Lo Presti, *Pentagon functions for massless planar scattering amplitudes*, *JHEP* **10** (2018) 103, [1807.09812].
- [282] D. Chicherin, T. Gehrmann, J. Henn, N. Lo Presti, V. Mitev and P. Wasser, *Analytic result for the nonplanar hexa-box integrals*, *JHEP* **03** (2019) 042, [1809.06240].
- [283] D. Chicherin, T. Gehrmann, J. Henn, P. Wasser, Y. Zhang and S. Zoia, *All Master Integrals for Three-Jet Production at Next-to-Next-to-Leading Order*, *Phys. Rev. Lett.* **123** (2019) 041603, [1812.11160].
- [284] C. Anastasiou and A. Lazopoulos, *Automatic integral reduction for higher order perturbative calculations*, *JHEP* **07** (2004) 046, [hep-ph/0404258].
- [285] C. Studerus, *Reduze-Feynman Integral Reduction in C++*, *Comput. Phys. Commun.* **181** (2010) 1293–1300, [0912.2546].
- [286] A. von Manteuffel and C. Studerus, *Reduze 2 - Distributed Feynman Integral Reduction*, 1201.4330.
- [287] R. N. Lee, *Presenting LiteRed: a tool for the Loop InTEgrals REDuction*, 1212.2685.
- [288] J. Gluza, K. Kajda and D. A. Kosower, *Towards a Basis for Planar Two-Loop Integrals*, *Phys. Rev. D* **83** (2011) 045012, [1009.0472].

- [289] R. M. Schabinger, *A New Algorithm For The Generation Of Unitarity-Compatible Integration By Parts Relations*, *JHEP* **01** (2012) 077, [1111.4220].
- [290] K. J. Larsen and Y. Zhang, *Integration-by-parts reductions from unitarity cuts and algebraic geometry*, *Phys. Rev. D* **93** (2016) 041701, [1511.01071].
- [291] J. Böhm, A. Georgoudis, K. J. Larsen, M. Schulze and Y. Zhang, *Complete sets of logarithmic vector fields for integration-by-parts identities of Feynman integrals*, *Phys. Rev. D* **98** (2018) 025023, [1712.09737].
- [292] J. Böhm, A. Georgoudis, K. J. Larsen, H. Schönemann and Y. Zhang, *Complete integration-by-parts reductions of the non-planar hexagon-box via module intersections*, *JHEP* **09** (2018) 024, [1805.01873].
- [293] J. Boehm, M. Wittmann, Z. Wu, Y. Xu and Y. Zhang, *IBP reduction coefficients made simple*, *JHEP* **12** (2020) 054, [2008.13194].
- [294] P. Mastrolia and S. Mizera, *Feynman Integrals and Intersection Theory*, *JHEP* **02** (2019) 139, [1810.03818].
- [295] H. Frellesvig, F. Gasparotto, M. K. Mandal, P. Mastrolia, L. Mattiazzi and S. Mizera, *Vector Space of Feynman Integrals and Multivariate Intersection Numbers*, *Phys. Rev. Lett.* **123** (2019) 201602, [1907.02000].
- [296] H. Frellesvig, F. Gasparotto, S. Laporta, M. K. Mandal, P. Mastrolia, L. Mattiazzi et al., *Decomposition of Feynman Integrals on the Maximal Cut by Intersection Numbers*, *JHEP* **05** (2019) 153, [1901.11510].
- [297] H. Frellesvig, F. Gasparotto, S. Laporta, M. K. Mandal, P. Mastrolia, L. Mattiazzi et al., *Decomposition of Feynman Integrals by Multivariate Intersection Numbers*, *JHEP* **03** (2021) 027, [2008.04823].
- [298] X. Liu, Y.-Q. Ma and C.-Y. Wang, *A Systematic and Efficient Method to Compute Multi-loop Master Integrals*, *Phys. Lett. B* **779** (2018) 353–357, [1711.09572].
- [299] X. Liu and Y.-Q. Ma, *Determining arbitrary Feynman integrals by vacuum integrals*, *Phys. Rev. D* **99** (2019) 071501, [1801.10523].
- [300] X. Guan, X. Liu and Y.-Q. Ma, *Complete reduction of integrals in two-loop five-light-parton scattering amplitudes*, *Chin. Phys. C* **44** (2020) 093106, [1912.09294].

- [301] P. Zhang, C.-Y. Wang, X. Liu, Y.-Q. Ma, C. Meng and K.-T. Chao, *Semi-analytical calculation of gluon fragmentation into $^1S_0^{[1,8]}$ quarkonia at next-to-leading order*, *JHEP* **04** (2019) 116, [1810.07656].
- [302] Y. Wang, Z. Li and N. Ul Basat, *Direct reduction of multiloop multiscale scattering amplitudes*, *Phys. Rev. D* **101** (2020) 076023, [1901.09390].
- [303] D. A. Kosower, *Direct Solution of Integration-by-Parts Systems*, *Phys. Rev. D* **98** (2018) 025008, [1804.00131].
- [304] G. De Laurentis and B. Page, *Ansätze for Scattering Amplitudes from p -adic Numbers and Algebraic Geometry*, 2203.04269.
- [305] L. Budge, J. M. Campbell, G. De Laurentis, R. K. Ellis and S. Seth, *The one-loop amplitudes for Higgs + 4 partons with full mass effects*, *JHEP* **05** (2020) 079, [2002.04018].
- [306] J. M. Campbell, G. De Laurentis, R. K. Ellis and S. Seth, *The $pp \rightarrow W(\rightarrow l\nu) + \gamma$ process at next-to-next-to-leading order*, *JHEP* **07** (2021) 079, [2105.00954].
- [307] J. M. Campbell, G. De Laurentis and R. K. Ellis, *Vector boson pair production at one loop: analytic results for the process $q\bar{q}\ell\bar{\ell}'\ell'g$* , 2203.17170.
- [308] S. Badger, C. Brønnum-Hansen, H. B. Hartanto and T. Peraro, *First look at two-loop five-gluon scattering in QCD*, *Phys. Rev. Lett.* **120** (2018) 092001, [1712.02229].
- [309] S. Abreu, F. Febres Cordero, H. Ita, B. Page and M. Zeng, *Planar Two-Loop Five-Gluon Amplitudes from Numerical Unitarity*, *Phys. Rev. D* **97** (2018) 116014, [1712.03946].
- [310] S. Badger, C. Brønnum-Hansen, T. Gehrmann, H. B. Hartanto, J. Henn, N. A. Lo Presti et al., *Applications of integrand reduction to two-loop five-point scattering amplitudes in QCD*, *PoS LL2018* (2018) 006, [1807.09709].
- [311] S. Abreu, F. Febres Cordero, H. Ita, B. Page and V. Sotnikov, *Planar Two-Loop Five-Parton Amplitudes from Numerical Unitarity*, *JHEP* **11** (2018) 116, [1809.09067].
- [312] T. Gehrmann, J. Henn and N. Lo Presti, *Analytic form of the two-loop planar five-gluon all-plus-helicity amplitude in QCD*, *Phys. Rev. Lett.* **116** (2016) 062001, [1511.05409].

- [313] S. Badger, C. Brønnum-Hansen, H. B. Hartanto and T. Peraro, *Analytic helicity amplitudes for two-loop five-gluon scattering: the single-minus case*, *JHEP* **01** (2019) 186, [1811.11699].
- [314] S. Abreu, J. Dormans, F. Febres Cordero, H. Ita and B. Page, *Analytic Form of Planar Two-Loop Five-Gluon Scattering Amplitudes in QCD*, *Phys. Rev. Lett.* **122** (2019) 082002, [1812.04586].
- [315] S. Abreu, J. Dormans, F. Febres Cordero, H. Ita, B. Page and V. Sotnikov, *Analytic Form of the Planar Two-Loop Five-Parton Scattering Amplitudes in QCD*, *JHEP* **05** (2019) 084, [1904.00945].
- [316] S. Abreu, F. F. Cordero, H. Ita, B. Page and V. Sotnikov, *Leading-color two-loop QCD corrections for three-jet production at hadron colliders*, *JHEP* **07** (2021) 095, [2102.13609].
- [317] S. Abreu, B. Page, E. Pascual and V. Sotnikov, *Leading-Color Two-Loop QCD Corrections for Three-Photon Production at Hadron Colliders*, *JHEP* **01** (2021) 078, [2010.15834].
- [318] H. A. Chawdhry, M. Czakon, A. Mitov and R. Poncelet, *Two-loop leading-color helicity amplitudes for three-photon production at the LHC*, *JHEP* **06** (2021) 150, [2012.13553].
- [319] H. A. Chawdhry, M. L. Czakon, A. Mitov and R. Poncelet, *NNLO QCD corrections to three-photon production at the LHC*, *JHEP* **02** (2020) 057, [1911.00479].
- [320] S. Kallweit, V. Sotnikov and M. Wiesemann, *Triphoton production at hadron colliders in NNLO QCD*, *Phys. Lett. B* **812** (2021) 136013, [2010.04681].
- [321] M. Czakon, A. Mitov and R. Poncelet, *Next-to-Next-to-Leading Order Study of Three-Jet Production at the LHC*, *Phys. Rev. Lett.* **127** (2021) 152001, [2106.05331].
- [322] S. Catani, *The Singular behavior of QCD amplitudes at two loop order*, *Phys. Lett. B* **427** (1998) 161–171, [hep-ph/9802439].
- [323] T. Becher and M. Neubert, *On the Structure of Infrared Singularities of Gauge-Theory Amplitudes*, *JHEP* **06** (2009) 081, [0903.1126].
- [324] T. Becher and M. Neubert, *Infrared singularities of scattering amplitudes in perturbative QCD*, *Phys. Rev. Lett.* **102** (2009) 162001, [0901.0722].

- [325] E. Gardi and L. Magnea, *Factorization constraints for soft anomalous dimensions in QCD scattering amplitudes*, *JHEP* **03** (2009) 079, [0901.1091].
- [326] E. Gardi and L. Magnea, *Infrared singularities in QCD amplitudes*, *Frascati Phys. Ser.* **50** (2010) 137–157, [0908.3273].
- [327] H. B. Hartanto, S. Badger, C. Brønnum-Hansen and T. Peraro, *A numerical evaluation of planar two-loop helicity amplitudes for a W-boson plus four partons*, *JHEP* **09** (2019) 119, [1906.11862].
- [328] S. Badger, E. Chaubey, H. B. Hartanto and R. Marzucca, *Two-loop leading colour QCD helicity amplitudes for top quark pair production in the gluon fusion channel*, *JHEP* **06** (2021) 163, [2102.13450].
- [329] P. Nogueira, *Automatic Feynman graph generation*, *J. Comput. Phys.* **105** (1993) 279–289.
- [330] J. Kuipers, T. Ueda, J. A. M. Vermaseren and J. Vollinga, *FORM version 4.0*, *Comput. Phys. Commun.* **184** (2013) 1453–1467, [1203.6543].
- [331] B. Ruijl, T. Ueda and J. Vermaseren, *FORM version 4.2*, 1707.06453.
- [332] G. Cullen, M. Koch-Janusz and T. Reiter, *Spinney: A Form Library for Helicity Spinors*, *Comput. Phys. Commun.* **182** (2011) 2368–2387, [1008.0803].
- [333] R. Eden, P. Landshoff, D. Olive and J. Polkinghorne, *The Analytic S-Matrix*. Cambridge University Press, 2002.
- [334] F. Cachazo, *Sharpening The Leading Singularity*, 0803.1988.
- [335] N. Arkani-Hamed, J. L. Bourjaily, F. Cachazo and J. Trnka, *Local Integrals for Planar Scattering Amplitudes*, *JHEP* **06** (2012) 125, [1012.6032].
- [336] D. C. Dunbar and W. B. Perkins, *Two-loop five-point all plus helicity Yang-Mills amplitude*, *Phys. Rev. D* **93** (2016) 085029, [1603.07514].
- [337] S. Badger, D. Chicherin, T. Gehrmann, G. Heinrich, J. Henn, T. Peraro et al., *Analytic form of the full two-loop five-gluon all-plus helicity amplitude*, *Phys. Rev. Lett.* **123** (2019) 071601, [1905.03733].
- [338] D. C. Dunbar, J. H. Godwin, W. B. Perkins and J. M. W. Strong, *Color Dressed Unitarity and Recursion for Yang-Mills Two-Loop All-Plus Amplitudes*, *Phys. Rev. D* **101** (2020) 016009, [1911.06547].

- [339] G. De Laurentis and D. Maître, *Two-Loop Five-Parton Leading-Colour Finite Remainders in the Spinor-Helicity Formalism*, *JHEP* **02** (2021) 016, [2010.14525].
- [340] M. Heller and A. von Manteuffel, *Multivariate Apart: Generalized partial fractions*, *Comput. Phys. Commun.* **271** (2022) 108174, [2101.08283].
- [341] L. E. K., *Factorization of rational functions of several variables into partial fractions*, *Izvestiya Vysshikh Uchebnykh Zavedenii. Matematika* **47** (1978) .
- [342] A. Raichev, *Leinartas's partial fraction decomposition*, 1206.4740.
- [343] Z. Bern, L. J. Dixon and D. A. Kosower, *New QCD results from string theory*, in *International Conference on Strings 93*, 5, 1993, hep-th/9311026.
- [344] G. Mahlon, *Multi - gluon helicity amplitudes involving a quark loop*, *Phys. Rev. D* **49** (1994) 4438–4453, [hep-ph/9312276].
- [345] Z. Bern, G. Chalmers, L. J. Dixon and D. A. Kosower, *One loop N gluon amplitudes with maximal helicity violation via collinear limits*, *Phys. Rev. Lett.* **72** (1994) 2134–2137, [hep-ph/9312333].
- [346] J. Henn, B. Power and S. Zoia, *Conformal Invariance of the One-Loop All-Plus Helicity Scattering Amplitudes*, *JHEP* **02** (2020) 019, [1911.12142].
- [347] V. M. Braun, G. P. Korchemsky and D. Müller, *The Uses of conformal symmetry in QCD*, *Prog. Part. Nucl. Phys.* **51** (2003) 311–398, [hep-ph/0306057].
- [348] D. C. Dunbar, G. R. Jehu and W. B. Perkins, *The two-loop n-point all-plus helicity amplitude*, *Phys. Rev. D* **93** (2016) 125006, [1604.06631].
- [349] D. C. Dunbar, G. R. Jehu and W. B. Perkins, *Two-loop six gluon all plus helicity amplitude*, *Phys. Rev. Lett.* **117** (2016) 061602, [1605.06351].
- [350] S. Badger, G. Mogull and T. Peraro, *Local integrands for two-loop all-plus Yang-Mills amplitudes*, *JHEP* **08** (2016) 063, [1606.02244].
- [351] D. C. Dunbar, J. H. Godwin, G. R. Jehu and W. B. Perkins, *Analytic all-plus-helicity gluon amplitudes in QCD*, *Phys. Rev. D* **96** (2017) 116013, [1710.10071].
- [352] D. C. Dunbar, W. B. Perkins and J. M. W. Strong, *n-point QCD two-loop amplitude*, *Phys. Rev. D* **101** (2020) 076001, [2001.11347].

- [353] A. R. Dalgleish, D. C. Dunbar, W. B. Perkins and J. M. W. Strong, *Full color two-loop six-gluon all-plus helicity amplitude*, *Phys. Rev. D* **101** (2020) 076024, [2003.00897].
- [354] D. Chicherin, V. Sotnikov and S. Zoia, *Pentagon Functions for One-Mass Planar Scattering Amplitudes*, 2110.10111.
- [355] C. G. Papadopoulos and C. Wever, *Internal Reduction method for computing Feynman Integrals*, *JHEP* **02** (2020) 112, [1910.06275].
- [356] S. Abreu, H. Ita, F. Moriello, B. Page, W. Tschernow and M. Zeng, *Two-Loop Integrals for Planar Five-Point One-Mass Processes*, *JHEP* **11** (2020) 117, [2005.04195].
- [357] D. D. Canko, C. G. Papadopoulos and N. Syrrakos, *Analytic representation of all planar two-loop five-point Master Integrals with one off-shell leg*, *JHEP* **01** (2021) 199, [2009.13917].
- [358] N. Syrrakos, *Pentagon integrals to arbitrary order in the dimensional regulator*, *JHEP* **06** (2021) 037, [2012.10635].
- [359] S. Badger, H. B. Hartanto and S. Zoia, *Two-Loop QCD Corrections to Wbb^- Production at Hadron Colliders*, *Phys. Rev. Lett.* **127** (2021) 012001, [2102.02516].
- [360] J. Kryś, *Recent progress for five-particle two-loop scattering amplitudes with an off-shell leg*, in *20th International Workshop on Advanced Computing and Analysis Techniques in Physics Research: AI Decoded - Towards Sustainable, Diverse, Performant and Effective Scientific Computing*, 2, 2022, 2202.06653.
- [361] S. Badger, H. B. Hartanto, J. Kryś and S. Zoia, *Two-loop leading colour helicity amplitudes for $W^\pm\gamma + j$ production at the LHC*, 2201.04075.
- [362] S. Abreu, H. Ita, B. Page and W. Tschernow, *Two-Loop Hexa-Box Integrals for Non-Planar Five-Point One-Mass Processes*, 2107.14180.
- [363] A. Kardos, C. G. Papadopoulos, A. V. Smirnov, N. Syrrakos and C. Wever, *Two-loop non-planar hexa-box integrals with one massive leg*, 2201.07509.
- [364] S. Badger, M. Becchetti, E. Chaubey, R. Marzucca and F. Sarandrea, *One-loop QCD helicity amplitudes for $pp \rightarrow t\bar{t}j$ to $O(\epsilon^2)$* , 2201.12188.
- [365] G. Bevilacqua, H. B. Hartanto, M. Kraus, J. Nasufi and M. Worek, *NLO QCD corrections to full off-shell production of $t\bar{t}Z$ including leptonic decays*, 2203.15688.

- [366] Z. Bern, L. J. Dixon and D. A. Kosower, *One loop corrections to five gluon amplitudes*, *Phys. Rev. Lett.* **70** (1993) 2677–2680, [[hep-ph/9302280](#)].
- [367] X. Chen, T. Gehrmann, E. W. N. Glover and M. Jaquier, *Precise QCD predictions for the production of Higgs + jet final states*, *Phys. Lett. B* **740** (2015) 147–150, [[1408.5325](#)].
- [368] X. Chen, J. Cruz-Martinez, T. Gehrmann, E. W. N. Glover and M. Jaquier, *NNLO QCD corrections to Higgs boson production at large transverse momentum*, *JHEP* **10** (2016) 066, [[1607.08817](#)].
- [369] F. Buccioni, S. Pozzorini and M. Zoller, *On-the-fly reduction of open loops*, *Eur. Phys. J. C* **78** (2018) 70, [[1710.11452](#)].
- [370] F. Buccioni, J.-N. Lang, J. M. Lindert, P. Maierhöfer, S. Pozzorini, H. Zhang et al., *OpenLoops 2*, *Eur. Phys. J. C* **79** (2019) 866, [[1907.13071](#)].
- [371] J.-N. Lang et al., “OTTER: On-The-fly TENSOR Reduction.”
https://gitlab.com/jean-nicolas_l/otftr.
- [372] A. Denner, S. Dittmaier and L. Hofer, *COLLIER - A fortran-library for one-loop integrals*, *PoS LL2014* (2014) 071, [[1407.0087](#)].
- [373] A. Denner, S. Dittmaier and L. Hofer, *Collier: a fortran-based Complex One-Loop Library in Extended Regularizations*, *Comput. Phys. Commun.* **212** (2017) 220–238, [[1604.06792](#)].
- [374] J. C. Collins and D. E. Soper, *Angular Distribution of Dileptons in High-Energy Hadron Collisions*, *Phys. Rev. D* **16** (1977) 2219.
- [375] E. Richter-Was and Z. Was, *Separating electroweak and strong interactions in Drell–Yan processes at LHC: leptons angular distributions and reference frames*, *Eur. Phys. J. C* **76** (2016) 473, [[1605.05450](#)].
- [376] V. Del Duca, F. Maltoni, Z. Nagy and Z. Trocsanyi, *QCD radiative corrections to prompt diphoton production in association with a jet at hadron colliders*, *JHEP* **04** (2003) 059, [[hep-ph/0303012](#)].
- [377] T. Gehrmann, N. Greiner and G. Heinrich, *Photon isolation effects at NLO in $\gamma\gamma$ + jet final states in hadronic collisions*, *JHEP* **06** (2013) 058, [[1303.0824](#)].
- [378] A. Djouadi, M. Spira and P. M. Zerwas, *Production of Higgs bosons in proton colliders: QCD corrections*, *Phys. Lett. B* **264** (1991) 440–446.

-
- [379] D. de Florian, M. Grazzini and Z. Kunszt, *Higgs production with large transverse momentum in hadronic collisions at next-to-leading order*, *Phys. Rev. Lett.* **82** (1999) 5209–5212, [[hep-ph/9902483](#)].
- [380] V. Ravindran, J. Smith and W. L. Van Neerven, *Next-to-leading order QCD corrections to differential distributions of Higgs boson production in hadron hadron collisions*, *Nucl. Phys. B* **634** (2002) 247–290, [[hep-ph/0201114](#)].
- [381] F. F. Cordero, A. von Manteuffel and T. Neumann, *Computational challenges for multi-loop collider phenomenology*, 4, 2022, [2204.04200](#).

Index

ΔR , 57

η , 57

N_c , 15

N_f , 22

s_{ij} , 23

tr_5 , 24

p_T , 57

Amplitude, 12

 MHV, 31

 UHV, 32

 colour-ordered, 16

 Eikonal, 61

 helicity, 31

 leg, 13

 loop, 13, 14, 35

 loop-induced, 51

 mass dimension, 13

 multiplicity, 13

 partial, 16

 primitive, 16

 propagator, 13

 splitting, 59

 tree, 13

 vertex, 13

Asymptotic freedom, 6

Boson, 3

W^\pm , 3

Z , 3

 mass, 132

 gauge, 3

 gluon, 3

 Higgs, 3

 photon, 3

 isolation, 57

 scalar, 3

 vector, 3

Colour, 5

 basis, 16

 fundamental, 17

 Casimir

 adjoint, 17

 fundamental, 17

 decomposition, 15

 Dynkin index, 16

 expansion, 19

 factor, 16

 generator, 5

 matrix, 19

 number of, 15

 structure constant, 5

 sum, 20

Correlation, 58

 colour, 62, 63

 spin, 61

Coupling

 electromagnetic, 4

 running, 6

 strong, 5

- Cross section, 9
 - hard, 9
- Current
 - eikonal, 63
 - off-shell, 32
- Diphoton-plus-jet production, 49, 54
 - gluon fusion, 51, 54
- Divergence, 14
 - infrared, 14, 56
 - collinear, 56, 59
 - soft, 56, 61
 - ultraviolet, 14
- Fermion, 3
 - QCD ghost, 5
 - flavour, 3
 - lepton, 3
 - charged, 3
 - neutrino, 3
 - quark, 3
 - number of, 22
- Gauge
 - axial, 6
 - light-cone, 6
 - light-like, 6
 - covariant, 5
 - R_ξ , 4, 5
 - Feynman, 5
 - Landau, 5
 - fixing, 4
 - invariance, 4, 14, 16
- Ghosts, 5
- Hadron, 6
 - baryon, 6
 - meson, 6
- Hadron collider, 1
- Hadronisation, 12
- Helicity, 22
- Identity
 - completeness, 30
 - Fierz, 17, 26, 28
 - Integration-by-parts, 43
 - Schouten, 28
 - Ward, 13
- Integral
 - master, 42
 - scalar, 38
 - scaleless, 14
- Irreducible scalar product, 41
- Jet, 12, 57
- Kinematic degrees of freedom, 23, 34
- Levi-Civita symbol, 24
- Little group, 27, 30
 - weight, 30
- Lorentz group, 25
- Matrix
 - element, 9
 - gamma, 4, 26
 - Gram, 24
 - determinant, 24
 - Pauli, 22
 - sigma, 26
- Metric tensor, 22
- Momentum, 25
 - conservation, 22, 29
 - reference, 29
 - twistor, 34
 - parametrisation, 98
 - variables, 34
- Monte Carlo event generator, 9
- On-shell, 22, 26
- Parton, 8
 - distribution function, 8
- Phase space, 9

- infrared limit, 56
 - double-unresolved, 64
 - single-unresolved, 58
- cut, 57
- degrees of freedom, 10
- Polarisation vector, 29
- Pole, 14, 58
 - spurious, 58
- Polylogarithm, 36
- Precision
 - QCD, 1
 - phenomenology, 1
- Pseudoscalar, 24
- Quantum
 - chromodynamics, 3, 5
 - β -function, 6
 - electrodynamics, 3
- Recursion, 32
 - BCFW, 33
 - Berends-Giele, 32
- Regularisation
 - dimensional, 14
 - 't Hooft-Veltman, 15
 - dimensional regulator, 14
 - spin dimension, 14
 - infrared subtraction, 66
 - scheme, 66
- Renormalisation, 15
- Resummation, 11
- Scale
 - dimension test, 111
 - factorisation, 8
 - equation, 9
 - renormalisation, 14
 - seven-point variation, 51
- Singularity, 14
- Spinor, 25
 - Dirac, 25
 - helicity formalism, 27
 - Weyl, 25
 - antiholomorphic, 25
 - bispinor, 26
 - holomorphic, 25
 - left-handed, 25
 - right-handed, 25
- Standard Model, 3
- Symmetry
 - crossing, 22
 - parity, 24
- Transcendentality, 36
- Unitarity, 12
 - cut, 36
 - maximal, 38
 - generalised, 38