

# Higgs Boson Pair Production in Vector Boson Fusion at NLO QCD in HEFT

Master's Thesis of

Jens Braun

at the Department of Physics  
Institute for Theoretical Physics (ITP)

First examiner: Prof. Dr. Gudrun Heinrich

Second examiner: PD Dr. Stefan Gieseke

First advisor: Prof. Dr. Gudrun Heinrich

Second advisor: Dr. Marius Höfer

20. November 2023 – 20. November 2024



---

*Higgs Boson Pair Production in Vector Boson Fusion at NLO QCD in HEFT (Master's Thesis)*

I declare that I have developed and written the enclosed thesis completely by myself. I have not used any other than the aids that I have mentioned. I have marked all parts of the thesis that I have included from referenced literature, either in their original wording or paraphrasing their contents. I have followed the by-laws to implement scientific integrity at KIT.

**Karlsruhe, 20. November 2024**

.....  
(Jens Braun)



# Abstract

For the upcoming high luminosity phase of the Large Hadron Collider, precise predictions for the production of two Higgs bosons are essential. Especially promising in this context are effective field theories, because they allow incorporating large classes of beyond Standard Model physics in the calculations with relatively little effort. Both, high precision predictions for the Standard Model and effective field theory calculations in vector boson fusion, are available. However, no next-to-leading order calculation in quantum chromodynamics of this process within an effective field theory framework is available yet.

In this work, the next-to-leading order quantum chromodynamics corrections to the production of two Higgs bosons in the vector boson fusion channel in the Higgs effective field theory (HEFT) are presented. As a first step towards this goal, the full set of leading-order operators in the Higgs effective field theory is reduced to the ones relevant to this process. Next, to make the numerical calculations more feasible, several improvements in the amplitude generation and kinematics calculations are implemented in the one loop provider GOSAM, significantly improving the program's overall runtime performance. Finally, GOSAM is used together with the Monte Carlo event generator WHIZARD to perform the numerical study of the vector boson fusion process, which includes a parameterization of the total cross section for arbitrary anomalous coupling values. Lastly, some differential distributions are presented for several values of the anomalous couplings.



# Zusammenfassung

Für die bevorstehende high-luminosity Phase des Large Hadron Colliders sind präzise Vorhersagen für die Produktion von zwei Higgs Bosonen unverzichtbar. Besonders vielversprechend sind in diesem Kontext effektive Feldtheorien, die es erlauben eine große Klasse von Theorien jenseits des Standardmodells mit verhältnismäßig wenig Aufwand in die Rechnungen miteinzubeziehen. Beides, präzise Vorhersagen im Standardmodell und Rechnungen in effektiven Feldtheorien für den Vector Boson Fusion Prozess, sind verfügbar. Rechnungen mit next-to-leading order Präzision in der Quantenchromodynamik innerhalb einer effektiven Feldtheorie Frameworks sind jedoch bisher nicht für diesen Prozess verfügbar.

In dieser Arbeit werden die next-to-leading order Korrekturen der Quantenchromodynamik für die Produktion von zwei Higgs Bosonen im Vector Boson Fusion Kanal im Kontext der Higgs effektiven Feldtheorie (HEFT) berechnet. Für diesen Zweck wird zuerst die volle Menge der leading-order Operatoren der Higgs effektiven Feldtheorie auf die für den Prozess relevanten reduziert. Um die folgenden Rechnungen praktikabler zu machen, werden auch einige Verbesserungen in der Erzeugung der Amplituden und der Berechnung der kinematischen Größen in dem One-Loop-Provider GOSAM implementiert, welche die Laufzeit des Programms deutlich reduzieren. Letztendlich wird GOSAM zusammen mit dem Monte Carlo Event Generator WHIZARD benutzt, um die numerische Studie des Vector Boson Fusion Prozesses durchzuführen, welche eine Parametrisierung des totalen Wirkungsquerschnitts für beliebige Werte der anomalen Kopplungen enthält. Letztendlich werden noch einige kinematische Verteilungen für verschiedene Werte der anomalen Kopplungen präsentiert.



# Contents

<b>Abstract</b>	<b>i</b>
<b>Zusammenfassung</b>	<b>iii</b>
<b>1 Introduction</b>	<b>1</b>
<b>2 Theoretical Background</b>	<b>3</b>
2.1 The Standard Model of Particle Physics . . . . .	3
2.2 Effective Field Theories . . . . .	6
2.2.1 Standard Model Effective Field Theory . . . . .	8
2.2.2 Higgs Effective Field Theory . . . . .	9
2.3 Calculation of Hadronic Cross Sections . . . . .	13
<b>3 Modernizing GoSam</b>	<b>17</b>
3.1 The Structure of a Calculation in GoSam . . . . .	17
3.2 Improving the Code Generation and Compile Times . . . . .	22
3.3 Improving the Runtime . . . . .	24
<b>4 Di-Higgs Production in Vector Boson Fusion</b>	<b>31</b>
4.1 Process Specification . . . . .	31
4.2 Validation of the Standard Model Calculation . . . . .	35
4.3 Parameterization of the Total Cross Section in Terms of the Anomalous Couplings . . . . .	36
4.4 Differential Results . . . . .	41
<b>5 Summary and Outlook</b>	<b>45</b>
<b>Bibliography</b>	<b>47</b>



# 1 Introduction

Since the discovery of the Higgs boson at the Large Hadron Collider (LHC) in 2012 [1, 2], the particle spectrum of the Standard Model (SM) is complete. One of the main objectives of the LHC physics program has been the precise measurements of its properties ever since. Today, many of the couplings involving a single Higgs boson are known with a precision of  $\sim 10\%$  [3]. At the same time, the couplings involving two or more Higgs bosons remain largely unexplored with the currently available experimental datasets. However, some of these couplings, e.g. the Higgs self-coupling, are essential to the structure of the electroweak symmetry breaking (EWSB) mechanism, which is why their measurement is one of the principal goals of the High Luminosity LHC (HL-LHC) physics program. These measurements can not only serve to verify the SM though, they can equally also reveal deviations from it. Since these couplings are currently only weakly constrained, they can incorporate a wide range of beyond Standard Model (BSM) physics.

The multi-Higgs couplings can best be probed in the production process of two Higgs bosons, to which four distinct production modes contribute sizeably [4]. Of these modes, mostly the dominant production process, gluon fusion, has recently received attention in terms of EFT studies (e.g. [5, 6, 7, 8]). For vector boson fusion (VBF), the second most important production process, some EFT studies have also been performed in e.g. [9, 10]. However, so far high precision predictions for this process are only available in the SM [11, 12, 12, 13, 14, 15], while the mentioned EFT calculations are only at leading order accuracy. The aim of this work is now to somewhat bridge this gap by providing the production of two Higgs bosons in the VBF channel at next-to-leading order (NLO) precision in quantum chromodynamics (QCD) with anomalous couplings given in terms of the Higgs Effective Field Theory (HEFT).

The thesis starts in chapter 2 with a review of the SM, which is then extended systematically within the Effective Field Theory (EFT) approach in the subsequent section. This especially includes the Lagrangian used in the later calculations, which is derived in section 2.2.2, followed by a sketch of the steps to calculate observables from the model in section 2.3. Chapter 3 is then dedicated to GOSAM, the program used to supply the scattering amplitude for the actual numerical study. Several changes in version 3 of GOSAM are described, which predominantly serve to improve the program's performance. Finally, the calculation of the production of two Higgs bosons in the VBF channel at NLO QCD in HEFT are presented in chapter 4.



# 2 Theoretical Background

In this chapter, a description of the various theoretical methods used in the later chapters of this work is given. First, the Standard Model of particle physics (SM) is reviewed in section 2.1, which also serves to give the sign conventions used in this thesis. Then, in section 2.2, the Effective Field Theory (EFT) approach to parameterize low energy deviations from the SM in a largely model independent way is described. This includes overviews of the two most common EFTs in current LHC physics, the Standard Model Effective Field Theory (SMEFT) in section 2.2.1 and the Higgs Effective Field Theory (HEFT) in section 2.2.2. Finally, section 2.3 gives a summary of how collider observables can actually be calculated from the previously given models.

## 2.1 The Standard Model of Particle Physics

The starting point for most particle physics calculations is the SM or extensions thereof, as is the case here. Therefore, as a first step, the SM is reviewed. This section follows the notations and sign conventions of [16].

The SM consists of two parts, Quantum Chromodynamics (QCD), describing the strong interaction, and the electroweak (EW) interaction. More formally, the SM is constructed as a quantum field theory (QFT) invariant under local transformations of the three gauge groups

$$SU(3)_c \times SU(2)_L \times U(1)_Y. \quad (2.1)$$

Here,  $SU(3)_c$  corresponds to QCD and  $SU(2)_L \times U(1)_Y$  represents the EW interaction. Gauge invariance under a general  $SU(N)$  gauge group with coupling constant  $g$  and generators  $T^a$  can then be insured by including a gauge field  $F_\mu^a$ , which is introduced into the covariant derivative as

$$D_\mu = \partial_\mu - igT^a F_\mu^a. \quad (2.2)$$

To write the kinetic energy of the gauge field, the field strength tensor

$$F_{\mu\nu}^a = \partial_\mu F_\nu^a - \partial_\nu F_\mu^a + gf^{abc} F_\mu^b F_\nu^c \quad (2.3)$$

is introduced, where the gauge group's structure constant is defined by  $[T^a, T^b] = if^{abc}T^c$ .

To complete the specification of the SM, its particle content has to be defined. There are three generations of fermions, indexed by  $i$ , each containing an up-type quark  $u_i$ , a

**Table 2.1:** Representations and charges under the SM gauge groups for the SM fermions and scalar bosons.

	$Q^i$	$u_R^i$	$d_R^i$	$L^i$	$e_R^i$	$\Phi$
$SU(3)_c$	<b>3</b>	<b>3</b>	<b>3</b>	<b>1</b>	<b>1</b>	<b>1</b>
$SU(2)_L$	<b>2</b>	<b>1</b>	<b>1</b>	<b>2</b>	<b>1</b>	<b>2</b>
$U(1)_Y$	$\frac{1}{6}$	$\frac{2}{3}$	$-\frac{1}{3}$	$-\frac{1}{2}$	$-1$	$\frac{1}{2}$

down-type quark  $d_i$ , a charged lepton  $e_i$  and a neutral lepton  $\nu_i$ . The quarks live in the fundamental representation of QCD's  $SU(3)_c$ , which means a quark carries a color charge with three possible values. The only other particle interacting with color charges is QCD's gauge boson itself, the gluon  $G_\mu^a$ , living in the adjoint representation with eight possible states. The situation becomes slightly more complicated with the electroweak interaction, since it is not parity-conserving. Let  $\psi_{L/R}$  denote left- and right-handed fermions. The left-handed particles are then grouped in  $SU(2)_L$  doublets

$$Q^i = \begin{pmatrix} u_L^i \\ d_L^i \end{pmatrix} \quad \text{and} \quad L^i = \begin{pmatrix} \nu_L^i \\ e_L^i \end{pmatrix}, \quad (2.4)$$

the right-handed field are  $SU(2)_L$  singlets  $u_R^i$ ,  $d_R^i$  and  $e_R^i$ . Finally, the fermions are charged under the  $U(1)_Y$ , with  $Y$  being the *hypercharge*. All these representations and charges are summarized in table 2.1. The  $SU(2)_L$  gauge bosons are denoted by  $W_\mu^I$ , the  $U(1)$  gauge boson by  $B_\mu$ .

So far, all fields are massless. Simply introducing mass terms is not possible here though, since a mass term for any field but the Higgs doublet would break the gauge symmetry. The existence of massive particles is dictated by nature however, requiring a more intricate mechanism to introduce the needed mass terms. This can be done through *spontaneous symmetry breaking*, which respects the gauge symmetry and only breaks it through the vacuum state. For this mechanism, a complex scalar  $SU(2)_L$  doublet  $\Phi$  is introduced, which is called the Higgs doublet. This field is then coupled to the remaining fields such, that they acquire a mass when the Higgs doublet acquires a non-zero vacuum expectation value (vev). All in all, the total SM Lagrangian before symmetry breaking reads

$$\begin{aligned} \mathcal{L}_{\text{SM}} = & -\frac{1}{4}G_{\mu\nu}^a G^{a,\mu\nu} - \frac{1}{4}W_{\mu\nu}^I W^{I,\mu\nu} - \frac{1}{4}B_{\mu\nu} B^{\mu\nu} \\ & + \sum_{\psi} \bar{\psi} D_\mu \gamma^\mu \psi - \left( \bar{Q} y^d \Phi d_R + \bar{Q} y^u \tilde{\Phi} u_R + \bar{L} y^e \Phi e_R + \text{h.c.} \right) \\ & + (D_\mu \Phi)^\dagger (D^\mu \Phi) - V(\Phi) \end{aligned} \quad (2.5)$$

where the sum over the generations is implicit,  $\tilde{\Phi} = i\sigma_2\Phi^*$  is the charge-conjugated Higgs-doublet and the  $\psi$ -sum runs over all fermions. The Lagrangian also includes the in general complex Yukawa matrices  $y^{u/d/e}$  and the Higgs potential, which reads

$$V(\Phi) = -\mu^2\Phi^\dagger\Phi + \lambda(\Phi^\dagger\Phi)^2 \quad (2.6)$$

in the SM.

**Spontaneous Symmetry Breaking** When  $\mu^2$  and  $\lambda$  are positive, the Higgs potential assumes its minimum at a non-zero value,

$$|\Phi_{\min}|^2 = \frac{\mu^2}{2\lambda} =: \frac{v^2}{2}. \quad (2.7)$$

At this point, the choice of the vacuum state is arbitrary among those fulfilling this relation. The symmetry breaking has to obey the observations made in nature however, especially the photon not acquiring a mass. According to Goldstone's theorem, one massless Goldstone boson will appear for each generator broken by the ground state. When those fields are coupled to gauge bosons, they appear as longitudinal degrees of freedom of these gauge bosons, which thus acquire a mass. Therefore, for the photon to stay massless, one generator must remain unbroken in the vacuum. If the ground state is chosen as  $\Phi_0 = (0 \ v)$ , it is invariant under

$$Q\Phi_0 = (\tau_3 + Y) \begin{pmatrix} 0 \\ v \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} 0 \\ v \end{pmatrix} = \begin{pmatrix} 0 \\ v \end{pmatrix} = \Phi_0 \quad (2.8)$$

where  $\tau_3 = \sigma_3/2$  is the third generator of  $SU(2)$ . The operator  $Q$  then represents the unbroken symmetry, which is the electromagnetic  $U(1)_{\text{em}}$ . It is most convenient to parameterize the full Higgs doublet as

$$\Phi = \frac{1}{\sqrt{2}} e^{i\sigma_i\phi_i/v} \begin{pmatrix} 0 \\ v+h \end{pmatrix}, \quad (2.9)$$

where  $\sigma_i$  are the Pauli matrices. This parameterization can be understood as starting with the field  $(0, v+h)$  and then applying an arbitrary local  $SU(2)$  gauge transformation on it, which is parameterized by the fields  $\phi_i$ . In this form, the  $\phi_i$  immediately drop from the Higgs potential, then reducing to

$$V(\Phi) = V(h) = v^2\lambda h^2 + \lambda v h^3 + \frac{\lambda}{4}h^4 + \text{const.} \quad (2.10)$$

This introduces a mass term for the previously massless field  $h$ , as well as some self-coupling terms. The parameterization (2.9) also allows to easily see the impact of the symmetry breaking on the remaining terms of the SM Lagrangian. The Lagrangian is constructed to be invariant under  $SU(2)_L \times U(1)_Y$ , therefore a gauge transformation can be performed on the gauge fields  $W^I$  and  $B$  to exactly remove the  $SU(2)$  transformation of (2.9) everywhere in the Lagrangian. Therefore, the all following considerations can

take place in the unitary gauge without loss of generality, where the Goldstone bosons are not present. The Higgs doublet's kinetic term reads

$$(D_\mu \Phi)^\dagger (D^\mu \Phi) = \frac{1}{2}(\partial_\mu h)(\partial^\mu h) + \frac{v^2}{4} \left( g^2 W_\mu^+ W^{-,\mu} + \frac{g^2 + g'^2}{2} Z_\mu Z^\mu \right) \left( 1 + \frac{h}{v} \right)^2 \quad (2.11)$$

after symmetry breaking, where

$$W_\mu^\pm = \frac{W_1 \mp iW_2}{\sqrt{2}} \quad \text{and} \quad \begin{pmatrix} Z_\mu \\ A_\mu \end{pmatrix} = \begin{pmatrix} \cos \theta_w & -\sin \theta_w \\ \sin \theta_w & \cos \theta_w \end{pmatrix} \begin{pmatrix} W_\mu^3 \\ B_\mu \end{pmatrix}. \quad (2.12)$$

The angle diagonalizing the mass matrix of  $W^3$  and  $B$  is the Weinberg angle  $\theta_w$ , which is defined as

$$\tan \theta_w = \frac{g'}{g}. \quad (2.13)$$

Equation (2.11) includes mass terms for the physical  $W^\pm$  and  $Z$  bosons, as well as couplings between two of these gauge bosons and one or two Higgs bosons. In summary, the generated boson masses read

$$M_h = \sqrt{2\lambda}v, \quad M_W = \frac{gv}{2}, \quad M_Z = \frac{\sqrt{g^2 + g'^2}}{2} = \frac{gv}{2 \cos \theta_w} = \frac{M_W}{\cos \theta}. \quad (2.14)$$

Finally, also the fermions acquire a mass term when equation (2.9) is inserted into the Yukawa terms of the SM Lagrangian. The Yukawa couplings are not diagonal in general, because of which the generated mass matrix is also not diagonal. The Yukawa matrices are not even necessarily diagonalizable, which is why a singular value decomposition with a subsequent field redefinition of the fermions must be used to diagonalize the mass terms. This field redefinition changes the Lagrangian however, introducing a remnant unitary matrix in the charged weak current, which is called *Cabibbo-Kobayashi-Maskawa* (CKM) matrix. Since it can be complex in general, it breaks the CP-invariance of the SM. After the diagonalization procedure, the fermion masses read

$$m_f = \frac{vy_f}{\sqrt{2}} \quad (2.15)$$

where  $y_f$  is the diagonal Yukawa coupling to the fermion  $f$ .

In the form of (2.5), the Lagrangian can not be used for perturbative calculations yet. This description includes unphysical degrees of freedom for the gauge bosons, which results in non-invertible equations for their propagators. Therefore, to perform calculations, gauge fixing and ghost terms need to be included in the Lagrangian as well, which then cancel the unphysical degrees of freedom. These are assumed to be included implicitly in the remaining thesis.

## 2.2 Effective Field Theories

In its near 50-year history, no significant numerical deviation from the SM has ever been observed conclusively. It is not without its deficiencies though, as there are some

internal puzzles in the SM, like the strong CP problem or the hierarchy problem. And in addition to the more subtle theoretical problems, there are several phenomena entirely undescribed by the SM, most notably gravity and the apparent existence of dark matter. The SM can therefore not be universally valid, its agreement with data at the currently probed energy scales is nonetheless excellent. It can therefore sensibly be used as the starting point to construct theories allowing slight deviations from the SM, which can be done in a largely model-independent way as an EFT. This section is largely based on [17].

Starting from a *top-down* approach, assume that the full UV theory is known. In order for the theory to allow predictions at arbitrarily large energies, it must be renormalizable. However, the behavior of the UV theory below a certain cutoff scale  $\Lambda$  can be approximated as a non-renormalizable theory. The effective low energy theory will contain all particles propagating below the cutoff scale, all other particles are *integrated out*, i.e. they are replaced by a generally infinite number of non-renormalizable operators describing their influence on the propagating fields. All appearing operators have to respect the symmetries of the UV theory, through which also all coefficients, called *Wilson coefficients*, are determined.

The general form of the Lagrangian of one such theory can also be constructed without restricting oneself to a specific UV theory though, with a *bottom-up* approach. A specific EFT is determined by the symmetries its operators have to obey and the fields able to appear, after which all possible operators up to a certain power can be written down. The Wilson coefficients can then be regarded as free parameters, which means a priori any UV theory obeying the specified symmetries can be approximated by the EFT. This approximate description of the UV theory holds below a certain scale, which is determined by the UV theory and the maximum power of considered EFT operators. When sufficient data is available, the Wilson coefficients can be determined by fitting predictions of the EFT to data. Since the Wilson coefficients are fully determined by the UV theory, the set of non-vanishing Wilson coefficients can be compared with different theories to determine which UV theory is not able to describe the observed data. This procedure therefore allows choosing viable theories without requiring to calculate predictions for each theory, but only the theory's Wilson coefficients.

Another critical aspect of the bottom-up approach is the number of considered operators. The number of possible operators will generally be large, but many of them are not independent. The Lagrangian is invariant under a variety of operations, which allows some operators to be mapped onto others. Two important tools for this are integration by parts and field redefinitions. Integration by parts allows shifting derivatives inside operators, producing one or more operators with the derivative at different positions, whereas field redefinitions allow absorbing some combinations of operators in others. Applying both can reduce the final number of operators significantly, producing a basis of the operators at the given order.

The EFT formalism has been applied to a large variety of problems over the years, with some well known examples being Fermi theory describing the low energy four-fermion approximation of the weak interaction and chiral perturbation theory describing the

interactions between hadrons as low energy limit of QCD. For current LHC physics, especially two EFTs are frequently used: SMEFT and HEFT.

### 2.2.1 Standard Model Effective Field Theory

The present LHC observations are well described by the SM, which is the foundation of SMEFT. It assumes the only propagating fields to be the SM fields, and requires the SM gauge symmetries to be obeyed by all operators. The EFT expansion is ordered in terms of the operators' canonical dimensions, which is composed of the dimensions of the appearing fields and differential operators. Since the Lagrangian must have dimension four, the dimension of all operators with  $d > 4$  has to be reduced to four by multiplying them with an energy scale  $\Lambda^{4-d}$ . This scale is conventionally called *new physics* scale  $\Lambda$ , since it controls the energy at which the EFT starts to noticeably deviate from the SM. In Summary, the SMEFT Lagrangian reads

$$\mathcal{L}_{\text{SMEFT}} = \mathcal{L}_{\text{SM}} + \sum_{d=5}^{\infty} \frac{1}{\Lambda^{d-4}} \sum_i C_i^{(d)} O_i^{(d)} \quad (2.16)$$

where  $i$  enumerates all operators  $O_i$  appearing at a certain operator dimension  $d$  and  $C_i$  are the Wilson coefficients.

The first BSM contributions in this Lagrangian appear at dimension 5, which contains only a single type of non-redundant operator (e.g. [18]). This operator violates lepton number conservation and generates Majorana mass terms for the neutrinos. Both potentially play an important role in extensions of the SM, e.g. generating the neutrino masses in order to explain neutrino oscillations. Compared to what is probed at the LHC though, these effects are generally associated with different physics. Therefore, the dimension five terms are usually dropped for LHC studies. Here, the considerations conventionally begin at dimension 6, which contains 2499 baryon and lepton number conserving operators in 59 classes [19]. With increasing dimension, the number of operators grows exponentially [20], already reaching 44 807 at dimension 8 (note that the counting is slightly different from before, as hermitian conjugates are counted separately here). Constraining a number of parameters this large to reasonable accuracy requires a massive amount of data, which is why often only the more controllable number of dimension 6 operators is used.

One important aspect of the Higgs sector in the SMEFT is the correlations between some couplings. Consider the purely bosonic sector, i.e. the couplings between the Higgs boson and the electroweak vector bosons. The Higgs doublet is in the fundamental representation of  $SU(2)_L$ , the electroweak vector bosons are in the adjoint representations of  $SU(2)_L$  and  $U(1)$  respectively. Therefore, in order to build a valid operator containing only these fields, the Higgs doublet always has to appear in the form  $\Phi^\dagger \Phi$ . In unitary gauge and after symmetry breaking, this equates to the expression  $(v + h)^2$ . It is therefore impossible to couple a single Higgs boson to a purely bosonic operator without also producing a di-Higgs coupling at the same time. Hence, every

single-Higgs vertex is correlated with the associated di-Higgs vertex in the purely bosonic sector of SMEFT. This fact plays an important role in the distinction of SMEFT and HEFT.

### 2.2.2 Higgs Effective Field Theory

The second EFT for LHC physics described here is the HEFT, which is also the theory used in the calculations of chapter 4. HEFT builds upon largely the same principles as SMEFT: it also enforces the SM's  $SU(3)_c \times SU(2)_L \times U(1)_Y$  gauge symmetry and particle content. It relaxes one of the SM's assumptions though, the organization of the weak Goldstones and the Higgs boson in a complex  $SU(2)_L$  doublet. Its specific form is motivated by two requirements, the first one being the introduction of the longitudinal polarizations of the weak vector bosons. This requires three Goldstones, since three degrees of freedom are required for the longitudinal polarizations. Only introducing these three Goldstones does not produce a fully consistent UV theory though, since the amplitude of diagrams with these Goldstones as external states grows with energy, eventually violating unitarity of the  $\mathcal{S}$ -matrix. This immediately leads to the second requirement, exact unitarity at all energies. This is ensured by also introducing the Higgs boson, which cancels the amplitude's piece growing with energy.

The form of the Goldstones cannot be altered, since it would at the same time impact the electroweak symmetry breaking. The existence of a scalar Higgs boson is also experimentally verified, the requirement of it ensuring exact unitarity can be relaxed though. The theory is constructed as an EFT, which is assumed to be valid only up to a certain energy anyway. Breaking unitarity above that scale does not matter in the context of the EFT then, merely the actual UV model has to repair the unitarity somewhere above the cutoff scale. Therefore, in the HEFT, the Goldstones and the Higgs are independent, leading to a more general coupling structure. To construct the HEFT Lagrangian, first separating the Goldstones and the Higgs in the SM Lagrangian is illustrative. In the parameterization of equation (2.9), the Higgs potential is already independent of the Goldstones. For the kinetic term, consider only the case without gauge fields for simplicity. Then, by introducing the unitary matrix  $\mathbf{U} = \exp(i\sigma_i\phi_i/v)$  and the unit vectors  $\hat{e}_{1,2}$  in the respective direction, the Higgs doublet can be written

$$\Phi = \frac{1}{\sqrt{2}}\mathbf{U}(h+v)\hat{e}_2. \quad (2.17)$$

Inserting this into the kinetic term yields

$$\begin{aligned} (\partial_\mu\Phi)^\dagger(\partial^\mu\Phi) &= \frac{1}{2}\hat{e}_2^\top((\partial_\mu\mathbf{U})(h+v) + \mathbf{U}(\partial_\mu h))^\dagger((\partial^\mu\mathbf{U})(h+v) + \mathbf{U}(\partial^\mu h))\hat{e}_2 \\ &= \frac{1}{2}(\partial_\mu h)(\partial^\mu h) + \frac{1}{2}(\hat{e}_2^\top((\partial_\mu\mathbf{U})^\dagger(\partial^\mu\mathbf{U}))\hat{e}_2)(h+v)^2 \\ &\quad + \frac{1}{2}(\hat{e}_2^\top((\partial_\mu\mathbf{U})^\dagger\mathbf{U} + \mathbf{U}^\dagger(\partial_\mu\mathbf{U}))\hat{e}_2)(\partial^\mu h)(h+v). \end{aligned} \quad (2.18)$$

The second and third term can further be simplified by evaluating the derivative explicitly,

$$\partial_\mu \mathbf{U} = \frac{i\sigma_i(\partial_\mu \phi_i)}{v} \mathbf{U}. \quad (2.19)$$

This immediately implies

$$(\partial_\mu \mathbf{U})^\dagger \mathbf{U} + \mathbf{U}^\dagger (\partial_\mu \mathbf{U}) = \frac{i(\partial_\mu \phi_i)}{v} (\mathbf{U}^\dagger \sigma_i \mathbf{U} - \mathbf{U}^\dagger \sigma_i \mathbf{U}) = 0. \quad (2.20)$$

The second term can be expressed in a more convenient form by using

$$(\partial_\mu \mathbf{U})^\dagger (\partial^\mu \mathbf{U}) = \frac{1}{v^2} (\partial_\mu \phi_i) (\partial^\mu \phi_j) (\mathbf{U}^\dagger \sigma_i \sigma_j \mathbf{U}) = \frac{(\partial_\mu \phi_i) (\partial^\mu \phi_i)}{v^2} \mathbb{1} \quad (2.21)$$

and therefore

$$\hat{e}_2^\top ((\partial_\mu \mathbf{U})^\dagger (\partial^\mu \mathbf{U})) \hat{e}_2 = \frac{1}{2} \langle (\partial_\mu \mathbf{U})^\dagger (\partial^\mu \mathbf{U}) \rangle \quad (2.22)$$

where  $\langle \dots \rangle$  denotes the  $SU(2)$  matrix trace. Thus, the final form of the kinetic term reads

$$\mathcal{L}_{h,\text{kin}} = \frac{1}{2} (\partial_\mu h) (\partial^\mu h) + \frac{v^2}{4} \langle (\partial_\mu \mathbf{U})^\dagger (\partial^\mu \mathbf{U}) \rangle \left(1 + \frac{h}{v}\right)^2. \quad (2.23)$$

The Higgs and the Goldstones are essentially separate in this form, the only remnant of the grouping as a doublet is the appearance of the Higgs only in the form  $(v + h)^2$ . A similar procedure can also be performed on the Yukawa couplings, after which the whole Lagrangian is essentially separated in terms of Goldstones and the Higgs.

For the HEFT, the appearance of  $\mathbf{U}$  is kept unchanged in order to also keep the structure of the symmetry breaking. The Higgs is now a gauge singlet though, which allows it to generally couple arbitrarily to each gauge invariant operator. With this change, the leading order HEFT Lagrangian reads [21]

$$\begin{aligned} \mathcal{L}_{\text{LO}} = & -\frac{1}{4} G_{\mu\nu}^a G^{a,\mu\nu} - \frac{1}{4} W_{\mu\nu}^I W^{I,\mu\nu} - \frac{1}{4} B_{\mu\nu} B^{\mu\nu} + \sum_\psi \bar{\psi} D_\mu \gamma^\mu \psi \\ & - \frac{v}{\sqrt{2}} (\bar{Q} \mathcal{Y}^d(h) \mathbf{U} \hat{e}_2 d_R + \bar{Q} \mathcal{Y}^u(h) \mathbf{U} \hat{e}_1 u_R + \bar{L} \mathcal{Y}^e(h) \mathbf{U} \hat{e}_2 e_R + \text{h.c.}) \\ & + \frac{1}{2} (\partial_\mu h)^\dagger (\partial^\mu h) + \frac{v^2}{4} \langle (D_\mu \mathbf{U})^\dagger (D^\mu \mathbf{U}) \rangle \mathcal{F}(h) - V(h) \end{aligned} \quad (2.24)$$

where

$$V(h) = v^4 \sum_{k=2}^{\infty} v_k \left(\frac{h}{v}\right)^k, \quad \mathcal{F}(h) = 1 + \sum_{k=1}^{\infty} f_k \left(\frac{h}{v}\right)^k, \quad \mathcal{Y}^X(h) = y^X + \sum_{k=1}^{\infty} y_k^X \left(\frac{h}{v}\right)^k \quad (2.25)$$

are arbitrary power series in  $h/v$  with the normal SM Yukawa matrices  $y^X$ . The Lagrangian is now given in terms of the fundamental object  $\mathbf{U}$ , on which the covariant derivative acts as

$$D_\mu \mathbf{U} = \partial_\mu \mathbf{U} - ig W_\mu^I \tau^I \mathbf{U} + ig' B_\mu \mathbf{U} \tau_3 \quad (2.26)$$

with  $\tau_I = \sigma_I/2$  being the generators of  $SU(2)$ . In principle, the kinetic terms in (2.24) could also be dressed with a function of the same form as  $\mathcal{F}(h)$ , these are not considered though. For the kinetic terms of the Higgs and the fermions, these functions can be absorbed into the existing functions  $\mathcal{F}(h)$  and  $\mathcal{Y}^X(h)$  of equation (2.24) by field redefinitions of the Higgs and fermions respectively. For the kinetic terms of the gauge bosons, the reason is more phenomenological. Their field strengths are assumed to not couple strongly to the Higgs, which is why these terms are discarded for the leading Lagrangian. They are included in the higher order Lagrangians again, though.

Generally, the organization of the HEFT Lagrangian into orders is less straight forward than in SMEFT. Since the Lagrangian is constructed in terms of adimensional objects  $\mathcal{F}(h)$  and  $\mathbf{U}$ , a scheme based on canonical dimension to determine which operators are to be included at which order is not appropriate here. Instead, a more complex counting-scheme has to be employed. Following [22], the EFT expansion is organized as a loop-expansion. This follows from assuming that the NLO operators contribute with at least the size of the one-loop contributions of the LO operators, which arises naturally in strongly coupled UV scenarios. This results in the identification  $\Lambda \approx 4\pi v$ , where  $\Lambda$  is the EFT cutoff scale, which can also be reformulated as a dimension-based approach by introducing the *chiral dimension*. It is assigned according to

$$[\partial_\mu]_c = 1, \quad [\phi]_c = [h]_c = 0 \quad [X_{\mu\nu}]_c = 1, \quad [\psi]_c = \frac{1}{2}, \quad [g]_c = [y]_c = 1. \quad (2.27)$$

Here,  $\phi$  is a generic Goldstone boson,  $X_{\mu\nu}$  a gauge field strength tensor,  $\psi$  a fermion,  $g$  a gauge coupling and  $y$  a Yukawa coupling. Applying this to (2.24), all appearing operators possess chiral dimension 2. In turn, this implies that all operators with chiral dimension 2 should be included in the LO Lagrangian. There is one other operator satisfying this requirement,

$$\mathcal{L}_\beta = \beta v^2 \langle \mathbf{U}^\dagger (D_\mu \mathbf{U}) \tau_3 \rangle \langle \mathbf{U}^\dagger (D^\mu \mathbf{U}) \tau_3 \rangle. \quad (2.28)$$

For phenomenological reasons, this operator is also included only at next-to-leading order though. In unitary gauge, this operator reduces to

$$\mathcal{L}_\beta \sim \langle (-igW_\mu^I \tau^I + ig' B_\mu \tau_3) \tau_3 \rangle \langle (-igW^{I,\mu} \tau^I + ig' B^\mu \tau_3) \tau_3 \rangle \sim Z_\mu Z^\mu \quad (2.29)$$

with  $\langle \tau^I \tau^J \rangle = \delta^{IJ}/2$  and  $Z_\mu \sim gW_\mu^3 - g' B_\mu$ . This operator therefore influences the  $Z$ -mass, but not the  $W$ -mass. This then manifests in a change of the  $\rho$  parameter,

$$\rho = \frac{M_W^2}{M_Z^2 \cos^2 \theta_w}, \quad (2.30)$$

which is equal to one in the leading order SM according to equation (2.14). The  $\rho$  parameter also receives contributions from higher order perturbative corrections in the SM. After this is accounted for, the value for purely BSM contributions is [3]

$$\rho_{\text{BSM}} = 1.000\,31(19). \quad (2.31)$$

For this reason, the parameter  $\beta$  of the operator (2.29) must be small, which is why it is assumed to be subleading to the other contributions of (2.24).

The calculation of chapter 4 includes the one-loop contributions in QCD. Therefore, in the loop-based counting scheme, the operators of the NLO Lagrangian should a priori contribute at the same order. However, the one-loop contributions themselves are also ordered in the coupling constants, which makes the electroweak one-loop contributions subleading to the QCD ones. Since the considered EFT operators for this process only act in the electroweak sector, the corresponding NLO operators acting in the same sector are also considered to be of the same order as the perturbative electroweak NLO contributions. They are hence also considered subleading to the QCD one-loop contributions and therefore not included in this study. Note that this might not actually be the case, since the NLO electroweak contributions are of a similar order as the QCD ones [11]. Nevertheless, they are not included to keep the scope of this calculation reasonable. In summary, only the Lagrangian (2.24) is considered in the remainder of this thesis. Additionally, no Yukawas appear in the VBF process, which is why they are also discarded. Finally, since two Higgs bosons are produced, only operators with a fixed number of Higgs bosons are relevant here. Therefore, the functions  $\mathcal{F}$  and  $V$  are truncated, yielding as only relevant non-SM operators

$$\mathcal{L} \supset \frac{v^2}{4} \langle (D_\mu \mathbf{U})^\dagger (D^\mu \mathbf{U}) \rangle \left( 1 + 2g_{\text{HVV}} \frac{h}{v} + g_{\text{HHVV}} \frac{h^2}{v^2} \right) - g_{\text{HHH}} \lambda v h^3. \quad (2.32)$$

The operators are normalized such, that in the SM  $g_{\text{HVV}} = g_{\text{HHVV}} = g_{\text{HHH}} = 1$ . In summary, the SM vertices are modified by:

- All vertices coupling two Goldstone,  $W$  or  $Z$  bosons to *one* Higgs boson are modified by a factor of  $g_{\text{HVV}}$ .
- All vertices coupling two Goldstone,  $W$  or  $Z$  bosons to *two* Higgs bosons are modified by a factor of  $g_{\text{HHVV}}$ .
- The triple-Higgs self coupling is modified by a factor of  $g_{\text{HHH}}$ .

Since this constitutes only a reweighing of SM contributions without new Lorentz structures, this description is equal to the  $\kappa$ -framework often used in experimental analyses. For an arbitrary coupling  $c$ , the  $\kappa$  value is defined as

$$\kappa_c = \frac{c}{c_{\text{SM}}} \quad (2.33)$$

where  $c_{\text{SM}}$  is the SM's prediction for the coupling  $c$ . Therefore, comparisons with experimental bounds are easily possible by identifying  $\kappa_\lambda = g_{\text{HHH}}$ ,  $\kappa_V = g_{\text{HVV}}$  and  $\kappa_{2V} = g_{\text{HHVV}}$ .

Comparing this picture with SMEFT, there is one important distinction in the allowed deviations from the SM. The SMEFT can also lead to deviations from the SM in all three couplings relevant here, but they are not independent. As described at the end

of the previous section, the Higgs boson can only appear in the form  $(v + h)^2$  in the SMEFT's bosonic sector. For the present couplings, this requires

$$g_{\text{HVV}}^{\text{SMEFT}} = g_{\text{HHVV}}^{\text{SMEFT}}, \quad (2.34)$$

which is the principal fact able to discriminate between HEFT and SMEFT in nature when considering the vector boson fusion process.

## 2.3 Calculation of Hadronic Cross Sections

So far, the relevant theories and their Lagrangians were described. However, there are still many steps between a theory given in terms of a Lagrangian and an actual measurement at a collider. What is usually performed are scattering experiments, in which the most straight forward observable is the counting rate  $\dot{N}$  of a scattering process. There are two kinds of contributions to this quantity, the first ones being collider-specific quantities, e.g. the number of particles per colliding bunch. The second kind are the purely fundamental physical ones, which is the part of actual interest. For this reason, the counting rate is separated as

$$\dot{N} = L \cdot \sigma, \quad (2.35)$$

where  $L$  is the luminosity, which contains all collider-specific information. The second part, the cross section  $\sigma$  contains the fundamental physics described by the model. The cross section for a hard scattering process  $i \rightarrow f$  is given by

$$\sigma(i \rightarrow f) = \frac{1}{F} \int |\mathcal{M}_{fi}|^2 d\Phi \quad (2.36)$$

where  $F$  is the incoming particle flux,  $\mathcal{M}_{fi}$  is the matrix element of the process and

$$d\Phi = (4\pi)^4 \delta(p_i - p_f) \prod_{k=1}^n \frac{1}{2E_k} \frac{d^3\vec{p}_k}{(2\pi)^3} \quad (2.37)$$

is the  $n$ -particle final state phase space. Here,  $p_{i/f}$  is the sum of the incoming/outgoing momenta and  $n$  is the number of outgoing particles. This expression for the cross section holds for incoming particles with known momenta, which leads to another complication for hadron colliders. There, the colliding particles are composite states, hence the incoming particles for the hard scattering process are constituents of the hadrons, not the hadrons themselves. Calculating the hard process then requires knowledge of the kind of the extracted partons and the momenta they carry. This can only be given in terms of a probability density function, in this context called parton distribution function (PDF), and not even this function can be calculated from first principles. In the low-energy context of the hadrons, the perturbative expansion of QCD does not converge, making the whole perturbative approach break down. Therefore, the PDFs have to be determined experimentally. The PDF is given as  $f_i(x, \mu)$ , where  $i$  is the

extracted parton,  $x$  is the fraction of momentum the parton carries and  $\mu$  is the scale at which the hadron is probed. The total hadronic cross section then reads

$$\sigma = \sum_{i,j} \int f_i(x_1, \mu) f_j(x_2, \mu) \sigma_{\text{part}}(i(x_1 p_1) + j(x_2 p_2) \rightarrow X; \mu) dx_1 dx_2 \quad (2.38)$$

where  $\sigma_{\text{part}}$  is given by equation 2.36 and  $i, j$  sum over all partons that can be extracted from the incoming hadrons.

When performing this calculation, two distinct computationally expensive tasks have to be performed: computing the matrix element  $\mathcal{M}_{f_i}$  and performing the high-dimensional phase space integration. In the setup used in the following chapters, both tasks are performed by separate programs. The matrix element is provided by a one-loop provider (OLP) program, which is discussed further in chapter 3. The remaining part of the calculation is handled by a Monte Carlo event generator, which handles integrating the phase space and generating events from it. They are called *Monte Carlo* event generators, because both of the aforementioned tasks are handled with algorithms based on random numbers. First, the total cross section must be calculated. For a  $2 \rightarrow n$  process, the dimension of the phase space scales as  $D = 3n - \text{const.}$  This leads to very high-dimensional integrations already at a moderate amount of outgoing particles, which limits the algorithms for reasonably fast integration greatly. Analytical integration can be discarded right away due to the complexity of a general scattering amplitude, leaving only numerical integration. For low-dimensional integrations, many efficient integration algorithms are available, but the situation becomes much more sparse for high-dimensional integrals. Here, the only algorithms computing the integral with reasonable accuracy and in reasonable time are based on MC integration. An arbitrary integral over an integration volume  $\Omega$  can be computed as

$$I = \int_{\Omega} f(x) d^d x = \lim_{N \rightarrow \infty} \frac{\Omega}{N} \sum_{k=1}^N f(x_k) \quad (2.39)$$

where the points  $x_k$  are sampled uniformly in the integration volume. The variance of the sum on the r.h.s. scales with  $\sigma_I^2 \sim 1/\sqrt{N}$ , independently of the integration volume's dimension. Because of this scaling MC integration performs better than other integration methods in many dimensions, which are affected much more strongly by the curse of dimensionality. Plain MC integration is still very inefficient though, the efficiency of the procedure can be improved significantly by not sampling the points  $x_k$  uniformly, however. Let  $x_k$  be sampled from a probability density function  $p(x)$  in  $\Omega$ , then

$$I = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=1}^N \frac{f(x_k)}{p(x_k)}. \quad (2.40)$$

The variance of this estimate vanishes when  $p(x) = f(x)/I$ , i.e. when the  $x_k$  are distributed according to  $f(x)$  itself. This would require already knowing  $I$  though, which is not obviously available yet. An approximation of this can still be used however, by sampling  $x_k$  according to a distribution that approximates  $f(x)$ . The variant of MC integration most often used in event generators is VEGAS [23], which approximates

the integrand through piecewise uniform distributions. This approximation is improved throughout the procedure, which is separated into distinct sampling steps. In each step, the function is sampled a number of times. These samples are then used to improve the approximation of  $f(x)$  for the next step. This way, statistics is accumulated while simultaneously improving the integration grid, which is what VEGAS calls the boundaries of the piecewise uniform distributions. Finally, after the integration is finished, the integration grid can also be used to improve the efficiency of sampling from differential cross sections.

These algorithms are implemented in WHIZARD [24], which is the MC event generator used in the study of chapter 4.



## 3 Modernizing GoSam

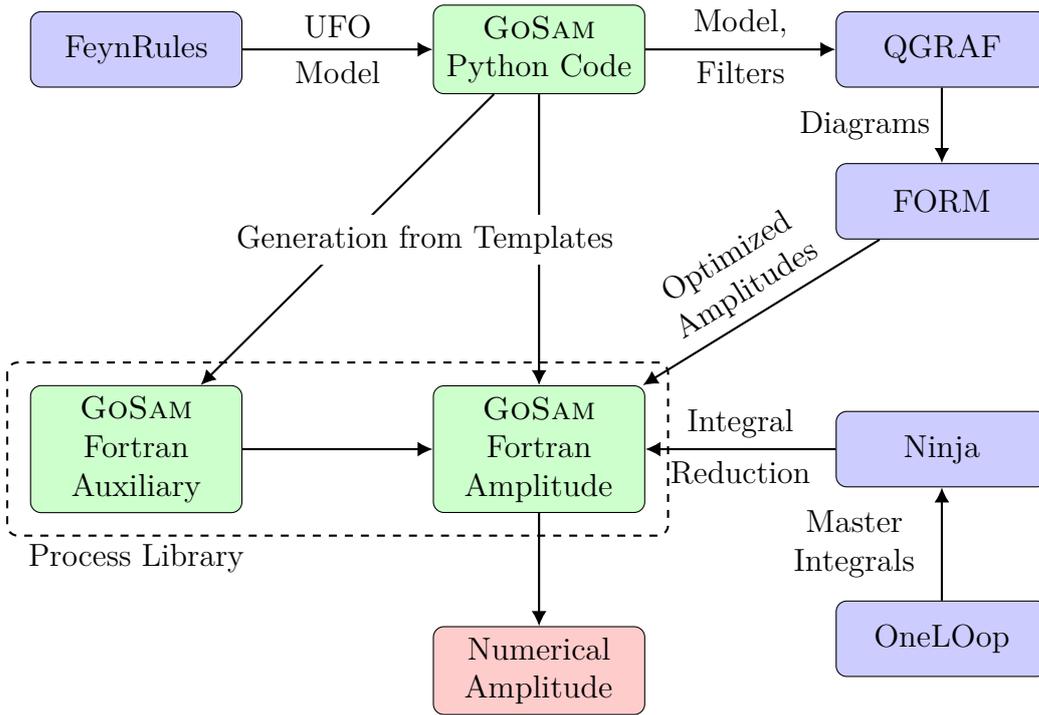
At the heart of a scattering calculation is the scattering amplitude, which is then further **process** by an MC event generator to determine total and differential cross sections. At the Born and one-loop level, many tools are available to largely automate the numerical evaluation of the scattering amplitude. These programs are generally referred to as one-loop providers, some available OLPs are e.g. GoSAM [25, 26], MADGRAPH5\_AMC@NLO [27], OPENLOOPS [28], RECOLA [29, 30] and VBFNLO [31, 32]. For this calculation, the OLP GoSAM is chosen because it generates the whole process code dynamically and thus places a priori no restrictions on the used model or the considered process. In practice, however, there are still technical restrictions like maximum tensor ranks of the appearing loop integrals limiting the choice of model and process. These are not the only problems limiting the usage of GoSAM for arbitrary calculations though, the largest obstacle is often simply the runtime of GoSAM and the generated process code.

This time obstacle is partly addressed in the new release GoSAM-3.0, which improves both the time GoSAM requires to generate the process library and the generated code's runtime performance. The changes made from GoSAM-2 to GoSAM-3 concerning performance are described in sections 3.2 and 3.3. Before the specifics of the new version are described though, the general workflow of GoSAM is reviewed in section 3.1.

### 3.1 The Structure of a Calculation in GoSam

To properly situate the internal changes in version 3 of GoSAM, this section sketches how a calculation is actually performed in GoSAM. The general workflow for GoSAM-3.0 is the same as in previous versions, which is described in [25, 26]. A schematic drawing of the whole workflow is depicted in figure 3.1.

As a first step, the physics of the theory has to be supplied. This is done in the form of a model containing most notably the particles, vertices and parameters. GoSAM itself ships with some built-in models covering the most frequently used models, e.g. the SM. Often, however, a more customized model is required. For this purpose, GoSAM supports the UFO-Format [33, 34]. This allows to supply an arbitrary model to GoSAM in a Python based format. There are also tools to automatically generate the UFO-Model from a physical Lagrangian, e.g. the Mathematica package FEYNRULES



**Figure 3.1:** Schematic structure of GoSAM’s calculation workflow. Green nodes are part of GoSAM or are generated by GoSAM, blue nodes are external programs.

[35], which is also what is commonly used alongside GoSAM. This is also the only component of figure 3.1 not included in the GoSAM installation itself.

With the model known, the GoSAM Python code processes the process specification. A process specification consists of one or more sets of incoming and outgoing particles together with the orders in the coupling constants. One such set of incoming and outgoing particles is called *subprocess*. The first simplification is already done at this stage. By considering only the external particles, crossing symmetry can be used to relate different subprocesses with each other. This allows some subprocesses to be expressed as another subprocess with crossed momenta, enabling GoSAM to reuse the code generated for one subprocess also for crossed subprocesses with minimal additional effort.

The next step is now to derive the amplitude expressions for each individual subprocess. GoSAM is a Feynman diagram based OLP, therefore the first step in the amplitude generation is to generate all possible Feynman diagrams contributing to the given subprocess. This step is delegated to an external program, QGRAF [36]. GoSAM translates the subprocess specification into a format suitable for QGRAF, which then generates the Feynman diagrams and outputs them in a format that can be further processed. Since the process specification only contains the coupling orders, QGRAF generally also generates all diagrams at those orders. This can be undesired when only a subset of the possible diagrams are to be taken into account. For this purpose, there are two possibilities to restrict the diagrams GoSAM processes further. The first

one are the filters inherent to QGRAF, which allow restricting e.g. the number of a specific vertex or a specific propagator in the generated diagrams. For even more sophisticated selection, QGRAF's output can be post-processed with Python filters. These allow to supply arbitrary Python code operating on a diagram, deciding whether the given diagram is to be kept or discarded. After the full set of Feynman diagrams for the subprocess is known, GOSAM performs some more simplifications on them, most notably diagram grouping. Not all diagrams are necessarily completely independent, some might e.g. share a common subdiagram. These shared components are identified and only processed once instead of for every diagram, significantly increasing the computational efficiency.

As the next step in the calculation, the Feynman diagrams have to be translated to their mathematical expressions. This task is performed with the symbolic manipulation toolkit FORM [37]. Using a specific style specification, QGRAF's output is a list of diagram objects built of external particles, vertices and propagators in FORM's syntax. Each of these components is then replaced with its mathematical representation by FORM. The resulting expressions are not usable for numerical evaluation yet though, first they need to be simplified and, in the case of one loop diagrams, the loop integration has to be carried out. Note that, for one-loop amplitudes, GoSam only constructs the numerator of the integrand, since this is the only part required for the integral reduction later on.

First, each diagram  $\mathcal{D}$  is decomposed in a process specific color basis  $|c_i\rangle$ ,

$$\mathcal{D} = \sum_i \mathcal{C}_i |c_i\rangle. \quad (3.1)$$

The  $|c_i\rangle$  consist only of products  $(T^{a_1} \dots T^{a_N})_{ij}$  of the generators and Kronecker deltas. To reach this form, the relations

$$iT_R f^{abc} = \text{Tr}(T^a T^b T^c) - \text{Tr}(T^a T^c T^b) \quad \text{and} \quad T_{ij}^a T_{kl}^a = T_R \left( \delta_{il} \delta_{kj} - \frac{1}{N_C} \delta_{ij} \delta_{kl} \right) \quad (3.2)$$

are used to remove all structure constants and color-contracted generators. Next, the partial amplitudes  $\mathcal{C}_i$  are further decomposed in the external helicities. The partial amplitudes can thus be calculated for given values of all external helicities and summed to get the full partial amplitude. One such partial amplitude with given external helicities is called *helicity amplitude*. With fixed helicities, the spinors of the massless fermions can be expressed in the convenient *spinor bracket* notation,

$$\begin{aligned} u_+(p) &= v_-(p) = |p\rangle, & u_-(p) &= v_+(p) = |p], \\ \bar{u}_+(p) &= \bar{v}_-(p) = \langle p|, & \bar{u}_-(p) &= \bar{v}_+(p) = [p|. \end{aligned} \quad (3.3)$$

The remaining wave functions are also expressed in terms of these spinor brackets. Massive fermions are written using a light cone decomposition,

$$p_i^\mu = l_i^\mu + \frac{p_i^2}{2p_i \cdot k} k^\mu \quad (3.4)$$

where  $k^\mu$  is an arbitrary lightlike reference vector. The spinors of massive fermions can then be expressed as

$$u_+(p) = |l\rangle + \frac{\sqrt{p^2}}{\langle pk\rangle} |k\rangle, \quad v_-(p) = |l\rangle - \frac{\sqrt{p^2}}{\langle pk\rangle} |k\rangle, \dots \quad (3.5)$$

with similar expressions for the other variations. Lastly, also the vector boson's polarization vectors are written in terms of the spinor brackets,

$$\varepsilon_\mu^+(p) = \frac{\langle q|\gamma_\mu|l\rangle}{\sqrt{2}\langle ql\rangle}, \quad \varepsilon_\mu^-(p) = \frac{[q|\gamma_\mu|l\rangle}{\sqrt{2}[ql]} \quad \text{and} \quad \varepsilon_\mu^0(p) = \frac{1}{\sqrt{p^2}}(2l_\mu - p_\mu), \quad (3.6)$$

where  $l = p$  in the massless case and the light cone decomposition of equation (3.4) otherwise. Additionally,  $q$  is an arbitrary reference vector and, since massless vector bosons cannot be longitudinally polarized,  $\varepsilon_\mu^0$  is only defined in the massive case.

At this point, the amplitude consists of constants, momenta, spinor brackets and Dirac matrices sandwiched between spinor brackets. Before the terms with the Dirac matrices can be further simplified, the regularization scheme has to be specified. GoSam uses dimensional regularization for both UV and IR divergences. More specifically, GoSam uses the dimensional reduction (DRED) scheme. In this scheme, the external vectors, i.e. momenta and polarization vectors, are kept four-dimensional and the Dirac algebra is also performed in four dimensions. Only the integration momentum is promoted to  $D = 4 - 2\varepsilon$ , but since the Dirac algebra is performed in four dimensions, this fact appears only at the end of the reduction. Then, the  $D$ -dimensional integration momentum  $q$  is split into a four-dimensional part  $\hat{q}$  and a remainder  $\tilde{q}$ . With  $\tilde{q}^2 =: -\mu^2$ , the simplifications can be fully performed in four dimensions and at the very end, the squared momentum is replaced by

$$q^2 \rightarrow \hat{q}^2 - \mu^2. \quad (3.7)$$

The actual Dirac algebra and further simplifications of the spinor bracket expressions are performed with the FORM library `spinney` [38]. This reduction is based on the Chisholm identities. Consider two strings of Dirac matrices  $\Gamma = \gamma^{\mu_1} \dots \gamma^{\mu_n}$  and  $\Gamma'$ , then

$$\langle p_i|\Gamma\gamma^\mu\Gamma'|p_j\rangle\gamma_\mu = 2\Gamma'|p_j\rangle\langle p_i|\Gamma - 2\tilde{\Gamma}|p_i\rangle\langle p_j|\tilde{\Gamma}' \quad (3.8)$$

where  $\tilde{\Gamma}$  is the same as  $\Gamma$  but with a reversed order of the Lorentz indices. Similar identities also exist for other combinations of angle and square spinor brackets. Finally, by applying the Chisholm identities together with

$$\gamma^\mu\Gamma\gamma_\mu = -2\tilde{\Gamma} \quad \text{and} \quad \gamma^\mu\Gamma\gamma^\nu\gamma_\mu = 2(\gamma^\nu\Gamma + \tilde{\Gamma}\gamma^\nu), \quad (3.9)$$

the spinor chains and traces of multiple Dirac matrices are expressed through terms containing fewer Dirac matrices. Lastly, contractions with Levi-Civita tensors are also written as combinations of spinor chains. After repeatedly applying all these identities, the amplitude is expressed through only a few fundamental objects: spinor

brackets like  $\langle p_i p_j \rangle$ ,  $[p_i p_j]$  and  $\langle p_i | \gamma^\mu | p_j \rangle$ , scalar products between four-dimensional momenta, constant parameters and the symbols  $\mu^2$  and  $\varepsilon$  from the dimensional reduction scheme.

For tree-level amplitudes, the end of the reduction is almost reached here. As a last step, the expressions for all diagrams are summed, once again simplified by FORM and finally written in a numerically more efficient form. Numerically more efficient here means that the expressions are rearranged with a Horner scheme to reduce the number of elemental numerical operations and subsequently a common subexpression elimination is applied to avoid evaluating the same subexpression multiple times. The final amplitude is then emitted as Fortran code through an appropriate template.

For one-loop diagrams, the amplitude consists of a loop integral over an optimized numerator at this stage. The actual loop integration is not handled by GOSAM, but rather by an integral reduction library interfaced to GOSAM. In version 1 of GOSAM, two integral libraries were available, GOLEM [39] and SAMURAI [40] (which is also the origin of the name GOSAM). In a later version, a third integral reduction library became available, NINJA [41]. GOLEM and SAMURAI are still available in GOSAM-3, but NINJA is generally faster and more stable compared to the two other options. For this reason, NINJA is generally preferred for calculations with GOSAM-3. Since it is the preferred option, only NINJA is sketched in the following.

NINJA is based on integrand reduction by Laurent expansion as described in [42]. With  $\mathcal{N}_0(\hat{q}, \mu^2)$  being the optimized numerator in DRED and  $r_i$  being a combination of external momenta, the general loop integral reads

$$\mathcal{A} = \int \frac{\mathcal{N}_0(\hat{q}, \mu^2)}{\prod_{i=0}^{n-1} D_i} d^D q \quad \text{with} \quad D_i = (\hat{q} - r_i)^2 - m_i^2 - \mu^2. \quad (3.10)$$

The method is based on now decomposing the loop integral in terms of a basis of scalar integrals, schematically

$$\mathcal{A} = \sum_{k=1}^4 \sum_{i_1 < \dots < i_k}^{n-1} \sum_j c_j^{(i_1 \dots i_k)} I_{i_1 \dots i_k} [f_j(\hat{q}, \mu^2)] \quad (3.11)$$

with

$$I_{i_1 \dots i_k} [f(\hat{q}, \mu^2)] = \int \frac{f(\hat{q}, \mu^2)}{D_{i_1} \dots D_{i_k}} d^D q. \quad (3.12)$$

The appearing scalar integrals are process independent and known analytically. They are provided to NINJA by a master integral library, in the case of GOSAM, ONELOOP [43] is used. The remaining task of NINJA is therefore to provide the process dependent coefficients  $c_j^{(i_1 \dots i_k)}$ . With the Laurent expansion method, only the numerator is required to be supplied. The coefficients are then extracted from the numerator by evaluating it at specific *unitarity-cuts*, i.e. at loop momenta where one or more propagators  $D_i$  vanish. The coefficient  $c_j^{(i_1 \dots i_k)}$  with  $k$  indices can be extracted from the numerator with  $k$  cut propagators, hence single-, double-, triple- and quadruple-cuts are required

to find all needed coefficients. For each such cut  $(i_1 \dots i_k)$ , there are solutions  $\hat{q}_\pm$  (only  $\hat{q}_+$  for the single-cut) of the loop momentum at which the numerator is evaluated. Then, schematically, the coefficients can be extracted through Laurent expansions of the general form

$$\left( \frac{\mathcal{N}_0(\hat{p}_\pm, \mu^2)}{\prod_{j \neq i_1 \dots i_k} D_j(\hat{p}_\pm, \mu^2)} - F \right) \Big|_{X \rightarrow \infty} = f(X; \{c_j^{(i_1 \dots i_k)}\}). \quad (3.13)$$

In the actually performed expansions, the expression  $F$  in the bracket can contain coefficients obtained by cutting more than  $k$  propagators, the variable  $X$  is defined by the parameterization of  $\hat{p}_\pm$  and is dependent on the considered cut. Finally,  $f$  is a Laurent series in  $X$  containing the  $c_j^{(i_1 \dots i_k)}$  as coefficients. By performing these expansions starting with the quadruple cut, all  $c_j^{(i_1 \dots i_k)}$  can be determined. NINJA provides the Laurent expansions through a numerical polynomial division for every phase space point. GOSAM therefore provides the optimized numerator as Fortran code in the format required by NINJA, which then computes the coefficients for the phase space point requested by the generated process library.

The reduction and simplification of the amplitude is complete at this point, but not all ingredients for the numerical evaluation are present. The amplitude is expressed through components like spinor brackets and scalar products, for which functions for numerical evaluation have to be provided as well. This is what is labeled as Fortran auxiliary functions in figure 3.1, which provide components like the kinematics, the color basis and some more utility functions. With the auxiliary part and the generated amplitude as the process library, linked to NINJA and ONELOOP, full numerical evaluation of tree-level and one-loop amplitudes is possible.

## 3.2 Improving the Code Generation and Compile Times

After execution of the GOSAM Python code concludes, its output consists of a process library skeleton containing the auxiliary functions, the final selection of diagrams and some FORM and Python code to handle the reduction and final Fortran code generation. The amount of generated Fortran code grows rapidly with the size of the considered process, easily reaching several tens of thousands of source files. This necessitates the use of a powerful build system to handle the generation and compilation of these many Fortran source files. Up to version 3, GOSAM provided two options for the build system: GNU Make [44] and GNU Autotools<sup>1</sup>. The actual code generation and compilation is specified in one or more Makefiles, which is Make's configuration format. These Makefiles could be generated in two ways, either directly by the GOSAM Python code or by Autotools. Autotools is a high level build system, which takes as input a general

<sup>1</sup>Autotools consists of several distinct programs like Automake and Autoconf, all of which can be found on <https://www.gnu.org/software>

configuration and a list of target files. It then generates a Makefile taking care of building the specified targets.

Make and Autotools are both very well-known and reliable build systems, although they clearly show their age in many places for very large projects like GOSAM's process libraries. A first version of Make was published in 1979, the different parts of Autotools were developed in the 1990s. Both dates are many years before the first multi-core processor was available, which is why both tools were originally designed to run serially on a single CPU core. Make was later equipped with the option to run multiple jobs in parallel, however, depending on the Makefile generated by Autotools, actually building a project in parallel can prove challenging. Additionally, the configuration for Autotools can quickly become cumbersome and verbose for large projects. For all of these reasons, GOSAM-3 instead uses a modern build system more suitable for building large project efficiently: `meson`<sup>2</sup>. `meson` takes the role of Autotools, being the high-level build system that generates the configuration for a low-level build system, which then takes care of actually building the targets. `meson`'s default low-level build system is `Ninja`<sup>3</sup> (not to be confused with the NINJA integral reduction library of section 3.1). Both are fundamentally designed to build large projects as fast as possible, and do so as automatically as possible. This includes, for example, automatically resolving dependencies between source files, which has to be done manually in Autotools projects.

In GOSAM-3, `meson` is used for two purposes: building the dynamically generated process library, and building GOSAM itself. The latter was previously done with an installation script `gosam_installer.py`, which used a pre-packaged version of the most of GOSAM's dependencies called `gosam-contrib`. In version 3, GOSAM's `meson` build definition directly downloads, builds and installs every dependency required to run a GOSAM calculation. This leads to a drastic reduction of the time required to install GOSAM: on an Intel Core i5-6500 4-core processor, installing GOSAM-3 takes roughly 150 s compared to roughly 400 s for GOSAM-2, even though both are configured to use all four cores. The picture is similar for `meson`'s second use in GOSAM, building the process library. GOSAM ships with a number of examples, each of which showcases some features of the program. There are 33 examples shared between GOSAM versions 2 and 3, providing a good overview of all available features. For this reason, the examples are also a good measure for performance differences between the GOSAM versions. For a given example, running `make test` executes the Python code, generates the process library, compiles everything and finally compares the squared amplitude to a reference value at a given phase space point. Computing the amplitude at a phase space point takes milliseconds in the worst case, hence the runtime of the whole procedure is completely dominated by the time it takes GOSAM to generate and compile the process library. The runtime of `make test` for most examples on an AMD Ryzen 7900X processor is shown in figure 3.2. Only the example `WpWpjj` is omitted, because it takes an order of magnitude longer to run compared to the other examples. Also note that

---

<sup>2</sup>The `meson` build system is available from <https://mesonbuild.com/>

<sup>3</sup>The `Ninja` build system is available from <https://ninja-build.org/>

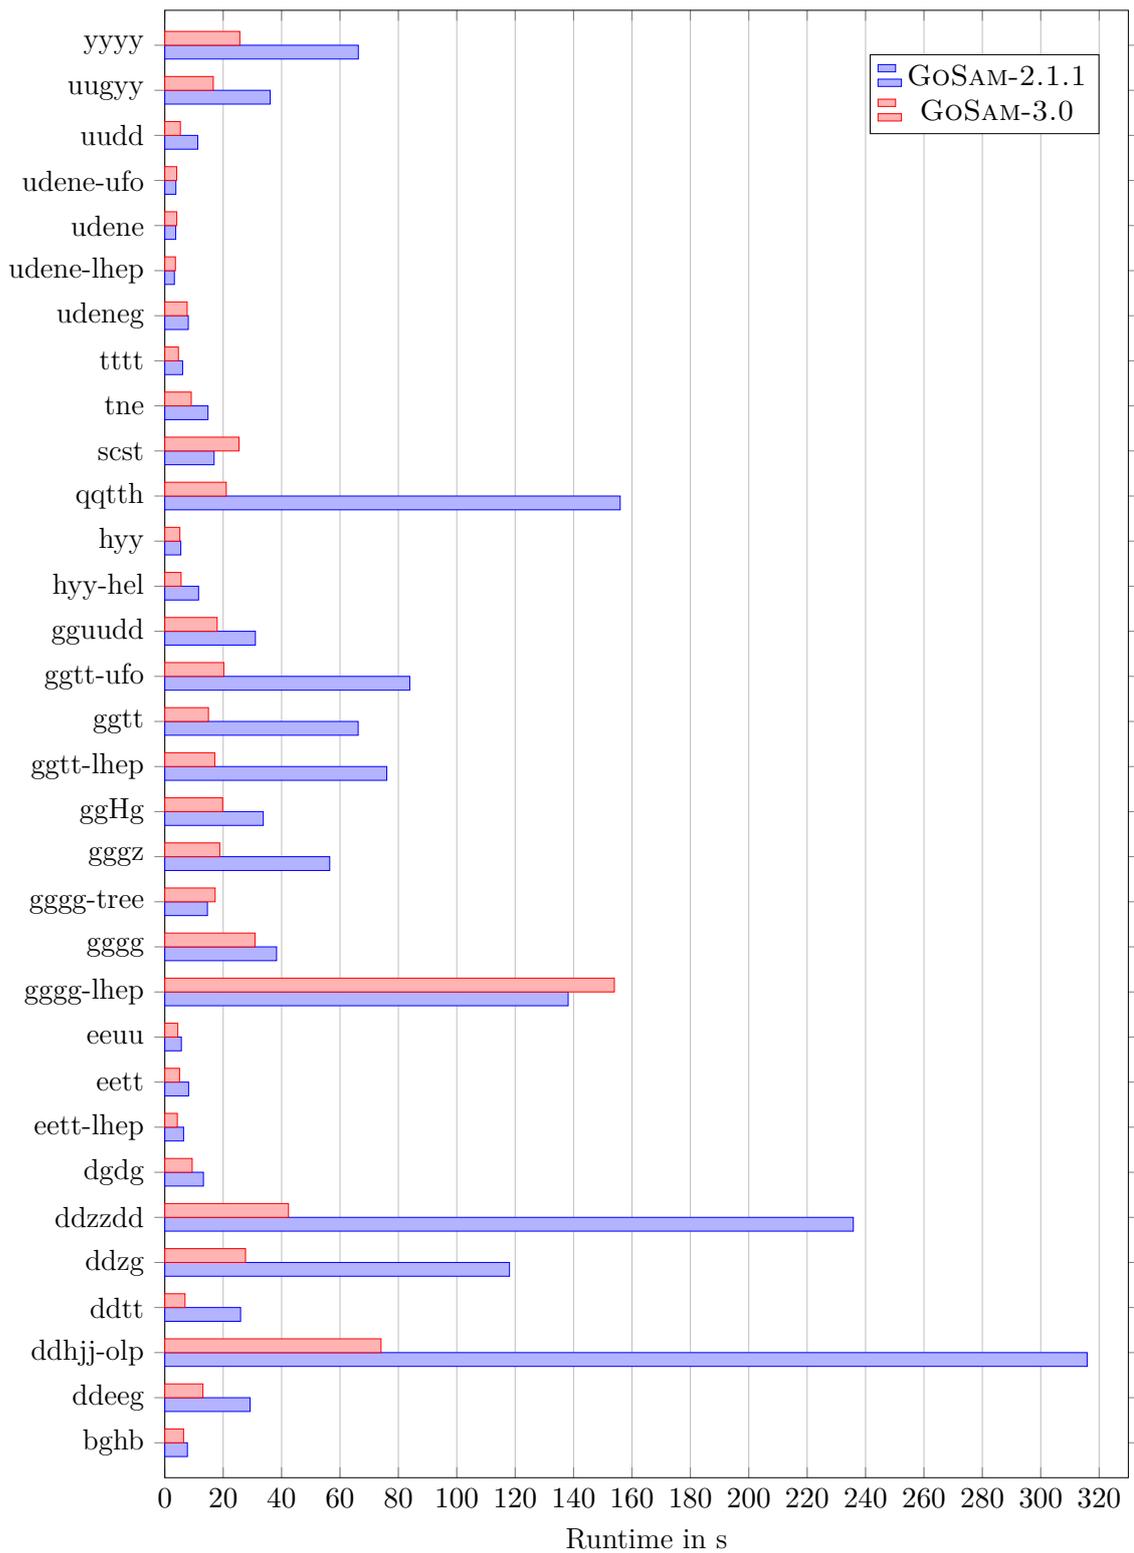
these measurements were not run in a strictly controlled environment and are therefore subject to runtime fluctuations due to operating system scheduling, processor clock speed fluctuations, etc. The runtimes in figure 3.2 are therefore only to be seen as estimates of what performance differences can be expected between GOSAM-2 and GOSAM-3, but not as rigorous performance metrics.

For the short running processes, e.g. `udene`, `hyy`, etc., the runtime difference is small, the reason being that these processes are also not fully parallelized by `meson`. For small processes, the runtime is mostly dictated by the runtime of the FORM reduction. Concerning the compilation, each source file can be processed simultaneously by `meson`, the code generation is only parallelized in the helicities and subprocesses though. This means that, for  $N_{\text{sp}}$  subprocesses and  $N_{\text{hel}}$  helicities per subprocess, the maximum number of utilized jobs during code generation is  $N_{\text{sp}}(N_{\text{hel}} + 1)$ . The additional one in the bracket is the contribution of the common auxiliary files (kinematics, color, etc.), which is also generated in parallel once per subprocess. Many of the examples are explicitly restricted to a single helicity in order to limit the runtime, at the same time also limiting the number of utilized jobs. In addition to the limited parallelization for some examples, the setup between the versions is also not completely identical. While the examples in GOSAM-2 only produce a static library, GOSAM-3 additionally generates a shared library, adding a step to the compilation and thus slightly increasing the runtime. In contrast to this, there are many examples showing significant improvements. This applies especially to the longer running examples, which benefit greatly from the parallelization and increased efficiency of `meson`. This is also the picture that generally holds for actual phenomenological studies, since these usually contain many subprocesses and are not restricted to specific helicities.

In summary, by transitioning the build infrastructure to the modern `meson` build system, building and installing the generated process library and GOSAM itself is more user-friendly and is able to much more efficiently saturate the capabilities of modern CPUs, significantly reducing the real runtime of a calculation in GOSAM-3.

### 3.3 Improving the Runtime

After the process library is generated and compiled, GOSAM's purpose concludes. The process library provides the squared scattering amplitude at a given phase space point, but calculating the (differential) cross section requires integrating the squared amplitude over the phase space. Since the phase space is high-dimensional for multi-leg processes, this can usually only be done by MC integration, requiring the squared amplitude to be evaluated millions of times for a reasonable estimate of the total cross section. Therefore, even if a single evaluation of the squared amplitude takes only milliseconds, the total time spent in the process library can quickly pile up to massive amounts during phase space integration. For this reason, making the evaluation of the squared amplitude as fast as possible is essential for large phenomenological studies.



**Figure 3.2:** Runtime of `make test` for all shared examples of GoSAM-3.0 and GoSAM-2.0, except `WpWpj`. Some example names are abbreviated to shrink the axis labels. A single run of `make test` includes running the Python code, generating and compiling the source and evaluation of a single phase space point.

Before describing changes specific to GOSAM, there are some purely technical changes in version 3 to improve the runtime. There are many optimizations available in the Fortran compiler, grouped into optimization levels. GOSAM-2 uses optimization level 2 by default, GOSAM-3 uses optimization level 3. This results in some additional compiler optimizations, which can slightly improve the runtime. As second technical change, GOSAM-3 allows choosing the target architecture<sup>4</sup>, which determines the set of CPU instructions used in the compiled program. Over time, many new instructions were added to CPUs, allowing some single instructions to perform tasks previously composited of multiple instructions. This reduces the required time for some tasks, but also makes older CPUs unable to execute the program correctly. Thus, the target architecture should always be chosen as the oldest out of all CPUs on which the program is supposed to be executed.

As mentioned in section 3.1, GOSAM already takes great care to generate efficiently evaluable amplitude code. With a Horner scheme reduction and a common subexpression evaluation, the number of elemental arithmetic operations is reduced. The amplitude code itself is not the only component contributing to the execution time though, the auxiliary functions also play an important role. To improve the runtime in GOSAM-3, changes have been made especially in two areas of the auxiliary functions, the model and the kinematics.

Looking at the model first, the principal change here is reducing the amount of duplicated code. Originally, GOSAM was designed to handle a single subprocess at a time, which is what GOSAM still does in standalone mode. Subsequently, it was equipped with an OLP mode, which allows simultaneous processing of multiple subprocesses through the BLHA interface [45, 46]. Because of this chronological order, the OLP mode is implemented such that GOSAM-2 treats each subprocess like in standalone mode, notably generating the auxiliary files once for each subprocess. This is required for some files like `color.f90` and `kinematics.f90`, since they are specific to the subprocess. Other files, however, are identical for all subprocesses and do not need to be generated multiple times. This applies also to `model.f90`, since all subprocesses in OLP mode are generated from the same model specification. Different subprocesses could be in general assigned different values of the parameters of an identical model, requiring each subprocess to have its own model. However, this is not intended in the BLHA2 interface, which only provides the `OLP_SetParameter` function to change parameters in all subprocesses simultaneously. Therefore, generating an identical `model.f90` for each subprocess is not required, but in itself also not a problem. Generating some identical files for each subprocess only impacts the compilation time slightly, the runtime of a calculation is not impacted at all. The runtime is only impaired in combination with changing a parameter value at each phase space point, e.g. when using running couplings. A model consists of external and internal parameters, values of external parameters are defined, the values of internal parameters are calculated from the external ones. When the values of an external parameter is changed, the internal

---

<sup>4</sup>A full list of possible target architectures for `gfortran` is available on <https://gcc.gnu.org/onlinedocs/gcc/x86-Options.html>

parameters dependent on the altered external parameter have to be recalculated. For simplicity, GOSAM does not check which internal parameters have to be recalculated and just recalculates all of them. In GOSAM-2, this happens for each subprocess at each phase space point. The time for setting a model parameter can then pile up, contributing a substantial part to a phase space integration for fast amplitudes. In GOSAM-3, some files, including `model.f90`, are shared between all subprocesses, thus reducing the time spent on changing the value of a model parameter by a factor given by the number of subprocesses.

Substantial changes are also made in the second aforementioned area, kinematics. After the reduction described in section 3.1 terminates, the amplitudes is given in terms of spinor brackets and scalar products. These are evaluated once per phase space point and can then be used multiple times when calculating the squared helicity amplitudes. In GOSAM, this is referred to as *kinematics* and is implemented in `kinematics.f90`. More specifically, three types of fundamental objects are calculated:

$$\langle pq \rangle = \text{Spaa}(p, q), \quad [pq] = \text{Spbb}(p, q) \text{ and } \langle p | \gamma_\mu | q \rangle = \text{Spab3\_vec}(p, q). \quad (3.14)$$

Here,  $p$  and  $q$  can be any lightlike external momenta with  $p \neq q$ . If the polarization vectors are evaluated numerically (the extension `numpolvec`, enabled by default), spinor brackets between polarization vectors and momenta are calculated as well. Before the helicity amplitude expressions are generated, it is a priori not known which spinor brackets will appear. For this reason, GOSAM-2 calculates all possible spinor brackets. Let  $n$  be the number of external particles, then the number of distinct pairs of external momenta is  $N = n(n-1)/2 \sim n^2$ . This is the number with which the amount of each type of spinor bracket scales, rapidly increasing with the number of external particles. Depending on the considered process, the actual number of required spinor brackets can be much smaller though. Because of this, GOSAM-2 can spend large amounts of time calculating spinor brackets that are never used, possibly substantially impairing the overall runtime. For this reason, GOSAM-3 takes care to only calculate the required spinor brackets. Like e.g. `color.f90`, `kinematics.f90` is now dynamically generated during the code generation step instead of from a static template by the Python code. This is done after the amplitude expressions are generated, allowing the script generating `kinematics.f90` to only include the spinor brackets actually appearing in the amplitude.

In addition to reducing the number of calculated spinor brackets, the actual time to calculate a spinor bracket of the kind  $\langle p | \gamma_\mu | q \rangle$  is reduced in GOSAM-3. In GOSAM-2, the bulk of this calculation happens in the function `Spab3_mcfm`, which, as the name suggests, is derived from the implementation in MCFM [47]. This function calculates

$$\text{Spab3\_mcfm}(p, k, q) = \langle p | \gamma_\mu | q \rangle k^\mu, \quad (3.15)$$

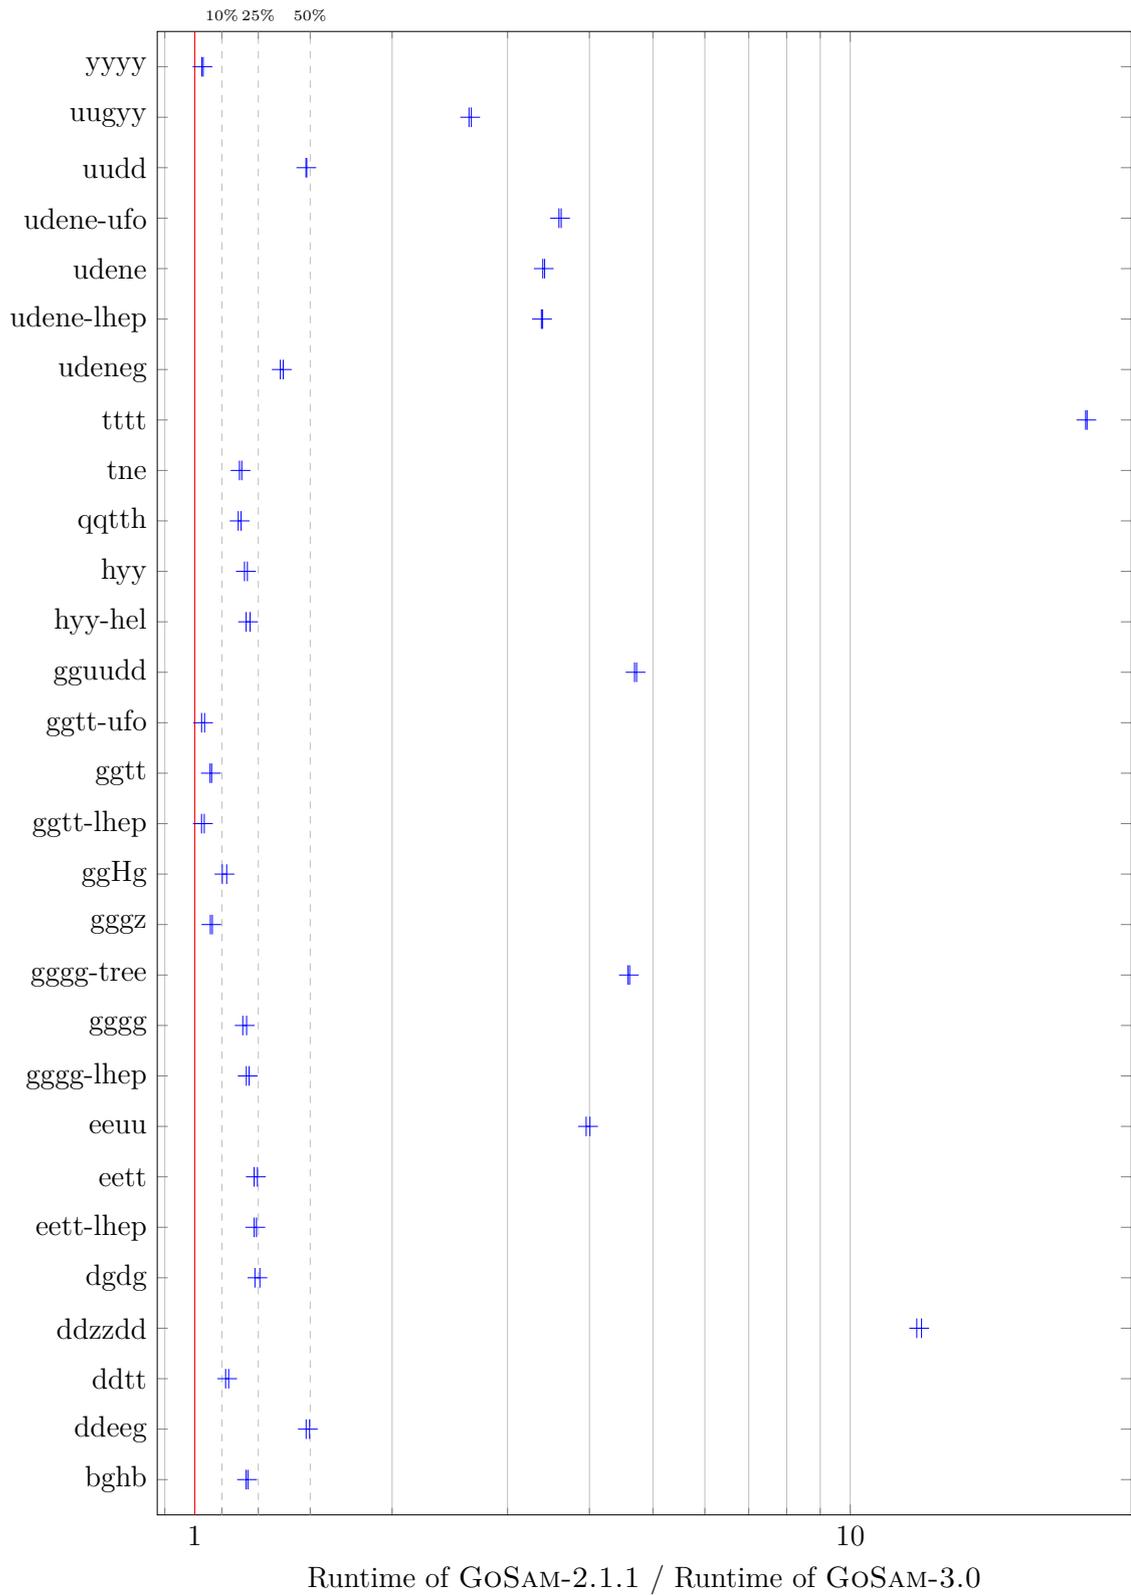
then

$$\text{Spab3\_vec}(p, q) = \langle p | \gamma_\mu | q \rangle = \langle p | \gamma_\nu | q \rangle (e_\mu)^\nu = \text{Spab3\_mcfm}(p, e_\mu, q) \quad (3.16)$$

where  $(e_\mu)^\nu = g_\mu^\nu$ . In this implementation, the function `Spab3_mcfm` is evaluated once for each spacetime component, unnecessarily recalculating intermediate results each time. This calculation is simplified in GOSAM-3, which calculates  $\langle p | \gamma_\mu | q \rangle$  directly. This reduces the runtime of a call to the function `Spab3_vec` from roughly 120 ns in GOSAM-2 to roughly 50 ns in GOSAM-3.

Lastly, the number of reevaluations of the whole kinematics per phase space point is reduced in GOSAM-3. GOSAM does not actually derive expressions for all possible helicities, but rather uses symmetries to map helicities onto each other. This generally amounts to permutations and reflections of momenta, which requires the kinematics to be reevaluated. Therefore, some helicity amplitudes share their kinematics, some helicity amplitudes have different kinematics. For this reason, the kinematics are reevaluated for each helicity amplitude in GOSAM-2, even if they are unchanged. The number of evaluations is reduced in GOSAM-3, in which the calculation of the helicity amplitudes is ordered such, that all helicities sharing the same kinematics are evaluated after each other. The kinematics are then only reevaluated once a helicity with a different kinematic configuration is reached. Depending on the considered process, this significantly reduces the time spent on kinematics.

The joint impact of all these changes can be seen in figure 3.3, which shows the runtime of GOSAM-2 relative to GOSAM-3's runtime for a call to `samplitude` across the examples. In GOSAM, `samplitude` is the function calculating the squared amplitude summed over the helicities at a given phase space point. The measurements are, again, taken on an AMD Ryzen 7900X processor and are only to be seen as estimates of possible differences. The runtime of all examples is improved in GOSAM-3, but the spread of runtime differences is large. The impact of the changes depends especially on two factors, the number of external particles and the complexity of the amplitude. The number of spinor brackets and helicities increases with the number external particles, which is why the changes generally become more effective with many external particles. At the same time though, the complexity of the amplitude increases, which also increases the time spent on calculating the actual amplitude. The relative impact of the kinematics on the total runtime thus decreases, making improvements in the kinematics on the total runtime less severe. This is also why the impact is generally lower for one-loop amplitudes, since their evaluation with NINJA generally takes much longer than evaluating a Born amplitude. This is very well represented in the examples, for which some processes like `yyyy` and `ggtt` show only very slight improvements. These are one-loop amplitudes which have a long total runtime and require almost all possible spinor brackets, keeping the impact of GOSAM-3's changes small. Most other one-loop examples see an improvement between 10% and 50%, with the bulk improving around 20%. However, there are also some outliers with larger differences among the one-loop examples, e.g. `udene` and `eeuu`. These are relatively simple processes with only few contributing diagrams, resulting in fast amplitudes where the kinematics' runtime significantly contribute. The most significant improvements can be seen in the tree-level processes though, `gguudd`, `gggg-tree`, `ddzzdd` and `tttt`. For them, GOSAM-2 requires several times the runtime of GOSAM-3 to evaluate the amplitude, reaching up to around 23 times longer for `tttt`.



**Figure 3.3:** Ratio of the runtime of a single call to `samplitude` for selected examples of GOSAM-3.0 and GOSAM-2.0. Some example names are abbreviated for smaller axis labels.

In spite of the large improvements in some amplitudes, the impact on runtimes of full phenomenological calculations is somewhat smaller. The largest runtime improvements are reached in Born amplitudes, which are generally much faster to evaluate than one-loop amplitudes, which obtain only smaller improvements. However, the improvements can have a large impact in the real corrections, which are tree-level amplitudes with an additional final state parton, contributing at the same order as the one-loop pieces. Since the dimension of their phase space is higher, they have to be evaluated many more times than the Born and one-loop contributions, which is why they can also contribute greatly to the total runtime of a phenomenological study. All in all though, the actual runtime difference between GOSAM-2 and GOSAM-3 significantly depends on the concrete process.

# 4 Di-Higgs Production in Vector Boson Fusion

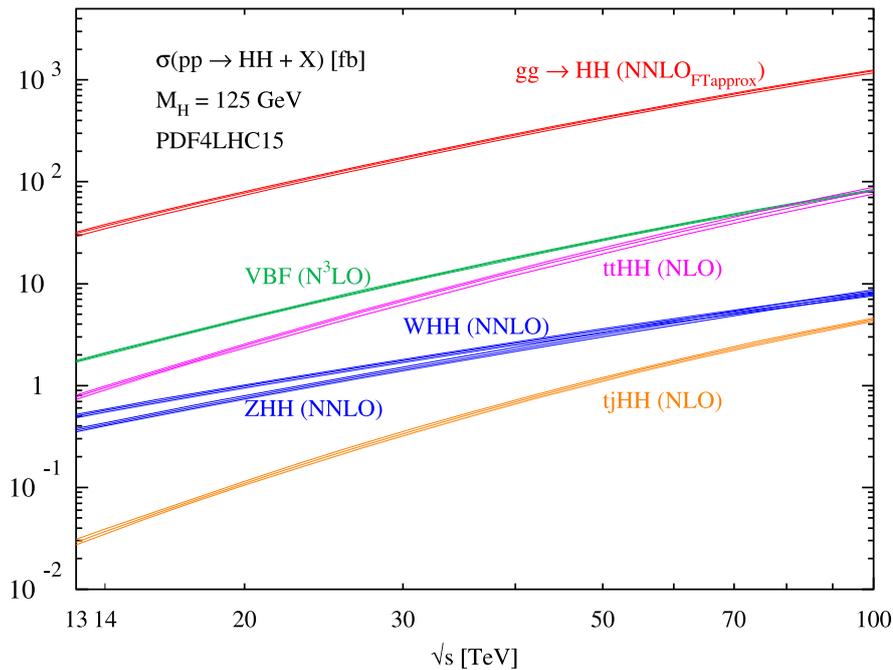
With the preliminaries of the numerical tools out of the way, this chapter will describe the actual phenomenological study of this work, the production of two Higgs bosons in the Vector Boson Fusion channel at the LHC. First, in section 4.1, the stage is set with a description of the studied process. Section 4.2 then continues with a general validation of the used simulation tool chain, which is compared to several other available SM tools. Finally, in the last two sections of this chapter, the results of this study are presented, starting with the total cross section in section 4.3 and closing with the differential distributions in section 4.4.

## 4.1 Process Specification

Similar to the single Higgs production case, there are four principal channels contributing to the production of two Higgs bosons at the LHC: Gluon Fusion (ggF), Vector Boson Fusion (VBF), double Higgs-Strahlung (VHH) and the associated production with top quarks (tHH) [4]. The production modes are listed in the order of their total cross sections, which are shown in figure 4.1 as a function of the total center of mass energy. The most important production channel is ggF, which has a cross section roughly an order of magnitude larger than all other channels combined at the LHC's center of mass energy. This makes ggF the prime channel to study in context of EFTs, as has been done in e.g. [5]. The ggF channel also comes with many complications though, especially it being loop-induced and therefore requiring two-loop amplitudes for NLO predictions. Compared to ggF, VBF has the obvious disadvantage of having a significantly smaller total cross section. Still, the VBF channel is very interesting to study, as it allows not only to probe the Higgs trilinear coupling  $g_{HHH}$ , but also the coupling of two vector bosons to two Higgs bosons,  $g_{HHVV}$ . And, since it also contains  $g_{HVV}$ , it allows testing for a possible correlation between  $g_{HVV}$  and  $g_{HHVV}$ , which would immediately discriminate between a realization of SMEFT or HEFT in nature.

Generally, the considered process in this study is the electroweak production of two Higgs bosons and two jets,

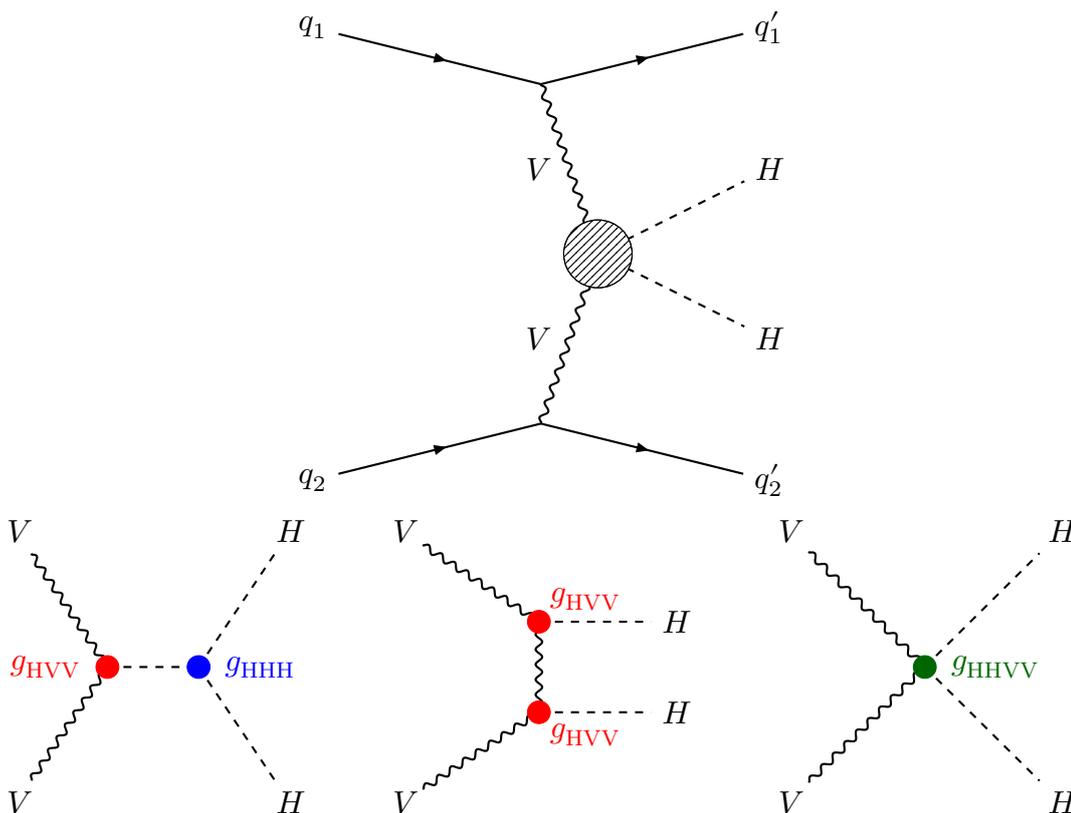
$$pp \rightarrow HHjj \quad \text{at} \quad \mathcal{O}(\alpha^4) \cdot \mathcal{O}(\alpha_S^{(0,1)}), \quad (4.1)$$



**Figure 4.1:** Total production cross sections of two Higgs bosons for the relevant channels as a function of the total center of mass energy. Figure taken from [4].

where  $\alpha_S^0$  is in the LO contributions and  $\alpha_S^1$  in the NLO contributions. The diagrams contributing at leading order are shown in figure 4.2. The upper diagram shows the general form of the  $t$ -channel contribution, the lower subdiagrams show the possible replacements for the gray blob in the upper diagram. In addition to the  $t$ -channel contribution, also  $u$ -channel and  $s$ -channel contributions are fully included. This is in contrast to the conventional definition of VBF, which does not include the  $s$ -channel contribution. They are in fact part of the double Higgs-Strahlung production mode, i.e. VHH where the produced vector boson decays to quarks. They are suppressed by the VBF cuts though, which require the jets to be strongly separated.

At next-to-leading order in QCD, a single gluon propagator is added for the virtual corrections. Like at leading order, no diagram selection is performed and all diagrams at the given orders of the coupling constants are included. This constitutes two classes of diagrams, the first one being the LO diagrams with a gluon propagator inserted to generate a loop. This includes vertex corrections with a triangle loop and diagrams with pentagon or hexagon loops, formed by connecting the two quark lines with the gluon. Examples of both are shown in the upper two diagrams in figure 4.3. The second class of diagrams are ones like the lower one in figure 4.3, which do not have a direct analog at leading order. Here, the quark lines are only connected by the gluon and the Higgs bosons are produced through a purely electroweak loop, which is only attached to one quark line. This also includes diagrams where the Higgses are produced in a bubble-like subdiagram, which is only attached to the quark line with a single propagator. Using e.g. FeynCalc [48], the total contribution of these diagrams can be seen to vanish analytically, which is why they are filtered away to improve GOSAM's



**Figure 4.2:** Feynman diagrams contributing to the VBF channel at leading order. The gray blob in the upper diagram can be substituted for any of the lower subdiagrams.

reduction time. For the real corrections, a single real gluon is simply attached to any of the four quark legs.

Model-wise, a modified version of the default FeynRules Standard Model<sup>1</sup> in Feynman gauge is used. This also includes the default settings, more specifically a diagonal CKM-Matrix and massless quarks. Additionally, **only first and second generation quarks are considered**. The model is modified in order to describe the HEFT Lagrangian, with the changes described in section 2.2.2. The input parameters are chosen to be

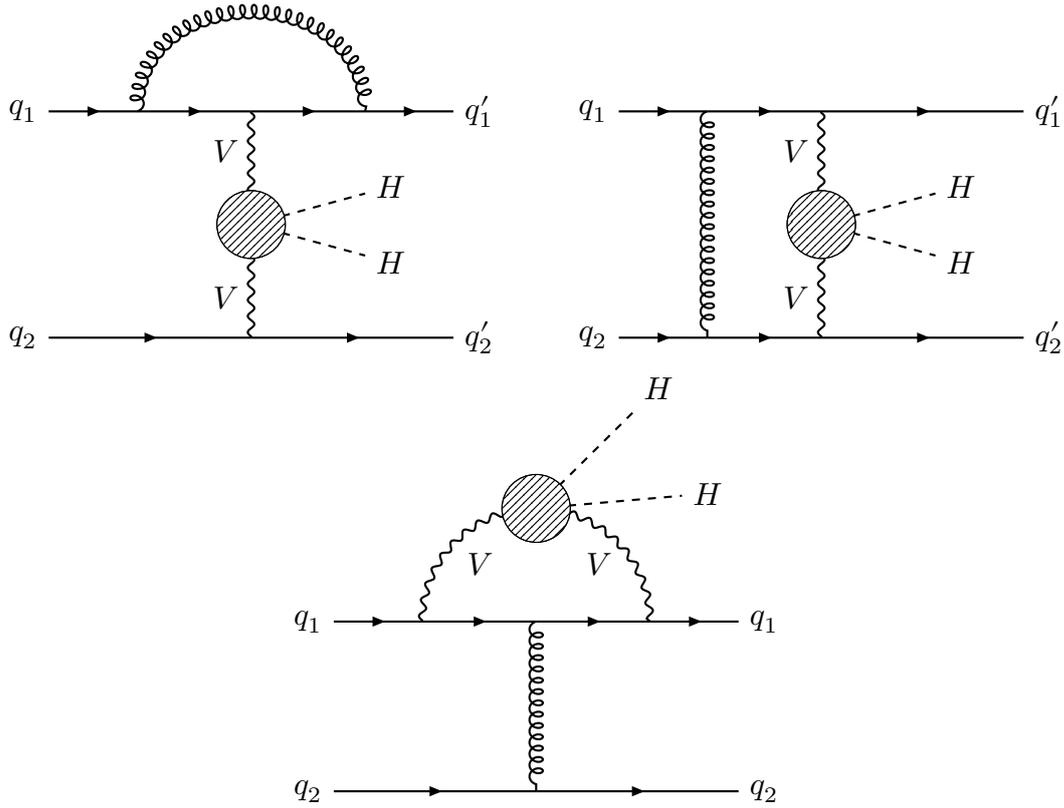
$$\frac{1}{\alpha(M_Z)} = 127.9, \quad G_F = 1.16637 \times 10^{-5}, \quad M_Z = 91.1876 \text{ GeV}, \quad M_H = 125.0 \text{ GeV}$$

$$\Gamma_Z = 2.4952 \text{ GeV}, \quad \Gamma_W = 2.085 \text{ GeV} \quad \text{and} \quad \Gamma_H = 4.07 \text{ MeV}. \quad (4.2)$$

All other parameters are derived from these, except for  $\alpha_S$ , which is taken from the PDF, for which PDF4LHC21\_mc (LHAPDF [49] ID 93000) is used. The central renormalization and factorization scale is chosen identically to [11],

$$\mu_0 = \sqrt{\frac{M_H}{2}} \sqrt{\frac{M_H^2}{4} + p_{\perp, HH}^2}, \quad (4.3)$$

<sup>1</sup><https://feynrules.irmp.ucl.ac.be/wiki/StandardModel>



**Figure 4.3:** Some diagrams contributing to the VBF channel at next-to-leading order in QCD.

where  $p_{\perp,HH}$  is the transverse momentum of the di-Higgs system. The scale uncertainties are then estimated by varying the renormalization and factorization scales with respect to this central scale. For this purpose a three-point scale variation scheme is employed, i.e.  $\mu_F = \mu_R = \xi\mu_0$  with  $\xi \in \{0.5, 1, 2\}$ .

Finally, some cuts are applied to the jets. They are clustered with the anti- $k_t$  algorithm [50] implemented in FastJet [51] and a radius parameter  $R = 0.4$ . Let  $j_1$  be the jet with the highest transverse momentum and  $j_2$  the one with the second-highest transverse momentum, together called *tagging jets*. First, some generic jet cuts are applied to both tagging jets:

$$p_{\perp,j} > 20 \text{ GeV} \quad \text{and} \quad |y_j| < 4.5. \quad (4.4)$$

Additionally, to separate the VBF contribution from e.g. QCD background or the VHH contribution, specific VBF cuts are applied,

$$m_{j_1 j_2} > 600 \text{ GeV} \quad \text{and} \quad \Delta\eta(j_1, j_2) > 4.0. \quad (4.5)$$

Concerning the values of the EFT parameters  $g_{HHH}$ ,  $g_{HVV}$  and  $g_{HHVV}$ , the considered ranges are chosen roughly according to current experimental bounds [52, 53],

$$g_{HHH} \in [-1, 6], \quad g_{HVV} \in [0.9, 1.1], \quad g_{HHVV} \in [0.5, 1.5]. \quad (4.6)$$

For the phenomenological study, points  $(g_{\text{HHH}}, g_{\text{HVV}}, g_{\text{HHVV}})$  are chosen to explore this EFT-space volume.

## 4.2 Validation of the Standard Model Calculation

As a first step, the validity of the matrix elements provided by GOSAM is checked. This is done by comparing the results produced by WHIZARD+GOSAM with other available tools. Since no calculation of this process in HEFT is available yet, the comparison is only done for the SM, i.e.  $g_{\text{HHH}} = g_{\text{HVV}} = g_{\text{HHVV}} = 1$ .

At leading order in the SM, several tools are available for the calculation of this process. WHIZARD+GOSAM is compared with a selection of them: MADGRAPH5\_AMC@NLO (MG5), VBFNLO and OPENLOOPS. MG5 and VBFNLO each provide their own MC integrator and event generator, OPENLOOPS is interfaced through WHIZARD. Each tool is initialized with the parameters given in 4.1, a different scale is chosen though. For better comparability, a static central scale is chosen,

$$\mu_0 = 2M_H. \quad (4.7)$$

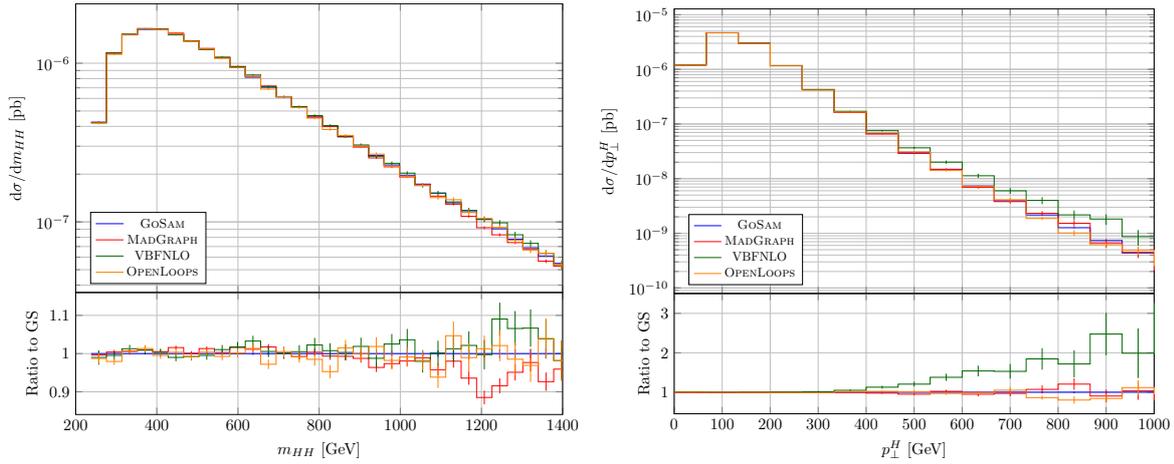
MG5 and OPENLOOPS use an identical process definition and also contain all possible diagrams at the relevant coupling order. This is not the case for VBFNLO, which implements the process only in the VBF approximation [54]. For the total cross sections, the tools yield

$$\begin{aligned} \sigma_{\text{GS}} &= 0.7165(3) \text{ fb}, & \sigma_{\text{OL}} &= 0.714(2) \text{ fb}, \\ \sigma_{\text{MG5}} &= 0.7146(3) \text{ fb}, & \sigma_{\text{VBFNLO}} &= 0.72128(8) \text{ fb} \end{aligned} \quad (4.8)$$

where GS represents GOSAM and OL represents OPENLOOPS. At high precision, there is a very slight disagreement between MG, GOSAM and VBFNLO. Since all of these tools use different MC event generators, the difference is most likely due to slight differences in e.g. technical cuts, or a different jet clustering algorithm in MG. VBFNLO deviates the most from the three other tools, which can also be seen in the  $p_{\perp}$  distributions of the Higgs bosons, depicted on the right side in figure 4.4. GOSAM, OPENLOOPS and MG show good agreement, while VBFNLO shows a clear excess in high- $p_{\perp}$ , which is unlikely to be solely due to statistical uncertainty. This, together with the slightly larger cross section, might be caused by the approximations made in VBFNLO and technical differences compared to the other tools. For all other considered distributions, all tools agree within the statistical uncertainties, e.g. in the  $m_{\text{HH}}$ -distribution depicted on the left side of figure 4.4.

At next-to-leading order, comparisons with MG5 and VBFNLO are not easily possible. MG5 cannot handle the appearing penta- and hexagon integrals, VBFNLO cannot generate NLO events on its own. For this reason, comparisons are only done with OPENLOOPS at NLO. The resulting total cross sections then read

$$\sigma_{\text{GS}} = 0.750(5) \text{ fb} \quad \text{and} \quad \sigma_{\text{OL}} = 0.743(6) \text{ fb}, \quad (4.9)$$



**Figure 4.4:** Distributions of the invariant mass of the Di-Higgs system and transverse momentum of all Higgs bosons calculated with GOSAM+WHIZARD, OPENLOOPS+WHIZARD, VBFNLO and MADGRAPH5\_AMC@NLO.

still being in good agreement at next-to-leading order.

In addition to comparing GOSAM and OPENLOOPS through (differential) cross sections, they are also directly compared on matrix element level. For this, a random phase space point is generated and the amplitude at this point is requested from GOSAM and OPENLOOPS. In the BLHA2 convention [45], this generally returns four values,  $A_2$ ,  $A_1$ ,  $A_0$  and  $T := |\mathcal{M}_{\text{Born}}|^2$ . For tree-level processes,  $A_2 = A_1 = A_0 = 0$ . For each value, the relative difference

$$\Delta_X = \frac{X_{\text{GS}} - X_{\text{OL}}}{X_{\text{GS}} + X_{\text{OL}}} \quad (4.10)$$

is calculated. With  $10^6$  points for each of the Born-, real- and loop-components, the orders of magnitude of the mean relative differences read

$$\begin{aligned} \Delta_T^{\text{Born}} &\sim 10^{-15}, & \Delta_T^{\text{Real}} &\sim 10^{-15}, \\ \Delta_{A_2}^{\text{Loop}} &\sim 10^{-14}, & \Delta_{A_1}^{\text{Loop}} &\sim 10^{-13}, & \Delta_{A_0}^{\text{Loop}} &\sim 10^{-8}. \end{aligned} \quad (4.11)$$

Therefore, the amplitudes provided by GOSAM and OPENLOOPS show excellent agreement for all components. Hence, in summary, all tools show good agreement with GOSAM, which is thus considered validated for this process.

### 4.3 Parameterization of the Total Cross Section in Terms of the Anomalous Couplings

The first type of result from the present calculation is the total cross section. It is calculated for twelve parameter points within the volume given at the end of section 4.1. The points are chosen based on a parameter scan performed at the leading order, with

**Table 4.1:** LO and NLO results for the total cross section at  $\sqrt{s} = 13.6$  TeV. The asymmetrical errors are calculated with a three-point scale variation, the symmetrical error is the MC integration error. For the ratio in the last column, only the MC error is given.

$g_{\text{HHH}}$	$g_{\text{HVV}}$	$g_{\text{HHVV}}$	$\sigma_{\text{NLO}}$ in fb	$\sigma_{\text{LO}}$ in fb	$\sigma_{\text{NLO}}/\sigma_{\text{LO}}$
1	1	1	$0.752^{+0.010}_{-0.031} \pm 0.009$	$0.814^{+0.096}_{-0.068} \pm 0.003$	0.923(11)
-1	0.9	1.5	$2.10^{+0.01}_{-0.05} \pm 0.01$	$2.290^{+0.308}_{-0.245} \pm 0.004$	0.916(6)
-1	1.05	1.3	$3.35^{+0.01}_{-0.10} \pm 0.03$	$3.70^{+0.34}_{-0.29} \pm 0.01$	0.907(8)
0	1	1	$1.89^{+0.01}_{-0.02} \pm 0.02$	$2.10^{+0.19}_{-0.17} \pm 0.01$	0.900(9)
1	0.9	1	$0.375^{+0.001}_{-0.012} \pm 0.004$	$0.408^{+0.050}_{-0.039} \pm 0.001$	0.921(10)
1	1	0.5	$4.38^{+0.01}_{-0.18} \pm 0.03$	$4.80^{+0.54}_{-0.46} \pm 0.01$	0.913(6)
1	1	1.5	$1.50^{+0.02}_{-0.03} \pm 0.01$	$1.660^{+0.218}_{-0.177} \pm 0.003$	0.902(6)
2	0.9	1.4	$3.81^{+0.06}_{-0.05} \pm 0.02$	$4.23^{+0.48}_{-0.40} \pm 0.01$	0.900(6)
2	1	1	$0.616^{+0.019}_{-0.028} \pm 0.007$	$0.683^{+0.069}_{-0.059} \pm 0.003$	0.902(11)
3	1.1	0.5	$4.32^{+0.01}_{-0.06} \pm 0.03$	$4.74^{+0.59}_{-0.49} \pm 0.01$	0.911(6)
4	0.95	0.5	$2.40^{+0.01}_{-0.06} \pm 0.01$	$2.62^{+0.28}_{-0.23} \pm 0.01$	0.914(8)
6	1.1	1	$9.83^{+0.01}_{-0.30} \pm 0.07$	$10.90^{+0.94}_{-0.84} \pm 0.03$	0.903(7)

the values chosen in table 4.1 exhibiting distinct shape deviations from the SM. Table 4.1 shows the LO and NLO cross sections at  $\sqrt{s} = 13.6$  TeV as well as the  $K$ -factor  $K = \sigma_{\text{NLO}}/\sigma_{\text{LO}}$ . As an immediate observation, the total cross section depends strongly on the couplings. The  $K$ -factor, on the other hand, is largely constant  $\mathcal{O}(-10\%)$  with respect to the anomalous couplings. The independence of the  $K$ -factor on the EFT couplings can be understood by considering the analytical structure of the amplitude. Schematically, at leading-order the amplitude can be decomposed as

$$\mathcal{M} \sim J_q^\mu J_{q'}^\nu V_{\mu\nu} \quad (4.12)$$

where  $J$  is a quark current and  $V$  is the tensor describing the  $VV \rightarrow HH$  subdiagram. The anomalous couplings only appear in the tensor  $V$ , while the currents  $J$  are independent of them. As argued in e.g. [55], the dominant NLO QCD contributions to the (single or double) VBF process are the vertex corrections like the upper left diagram in figure 4.3. These only appear in the currents  $J$ , i.e. the NLO QCD corrections and the differences due to the anomalous couplings are essentially factorized in the amplitude. For this reason the EFT dependence mostly cancels in the  $K$ -factor, leaving only the contribution from the NLO QCD corrections independently of the coupling values.

Also, a clear reduction of the scale dependence compared to the LO values can be observed in table 4.1. The scale variation changes the LO cross sections of order  $\sim 10\%$ , while the NLO cross sections only change on the order of  $\sim 3\%$ .

**Table 4.2:** Fit results for the coefficients of equation (4.14) at  $\sqrt{s} = 13.6$  TeV and with the central scale given by equation (4.3). The uncertainties are those obtained in the fitting procedure.

Parameter	$\mu_F = \mu_r = \mu_0/2$	$\mu_F = \mu_r = \mu_0$	$\mu_F = \mu_r = 2\mu_0$
$A_0$	0.7011(46)	0.6889(35)	0.6830(30)
$A_1$	22.15(12)	21.71(9)	21.55(8)
$A_2$	11.86(7)	11.59(5)	11.55(4)
$A_3$	-6.139(42)	-6.025(33)	-5.984(27)
$A_4$	3.865(29)	3.786(23)	3.773(19)
$A_5$	-31.44(18)	-30.75(14)	-30.56(11)

To generally describe the dependence of the total cross section on the anomalous couplings, consider the diagrams in figures 4.2. As indicated in the lower subdiagrams, the different diagrams have distinct dependencies on the anomalous couplings. The amplitude can therefore be parameterized as

$$\mathcal{M} = \alpha g_{\text{HVV}} g_{\text{HHH}} + \beta g_{\text{HVV}}^2 + \gamma g_{\text{HHV}} \quad (4.13)$$

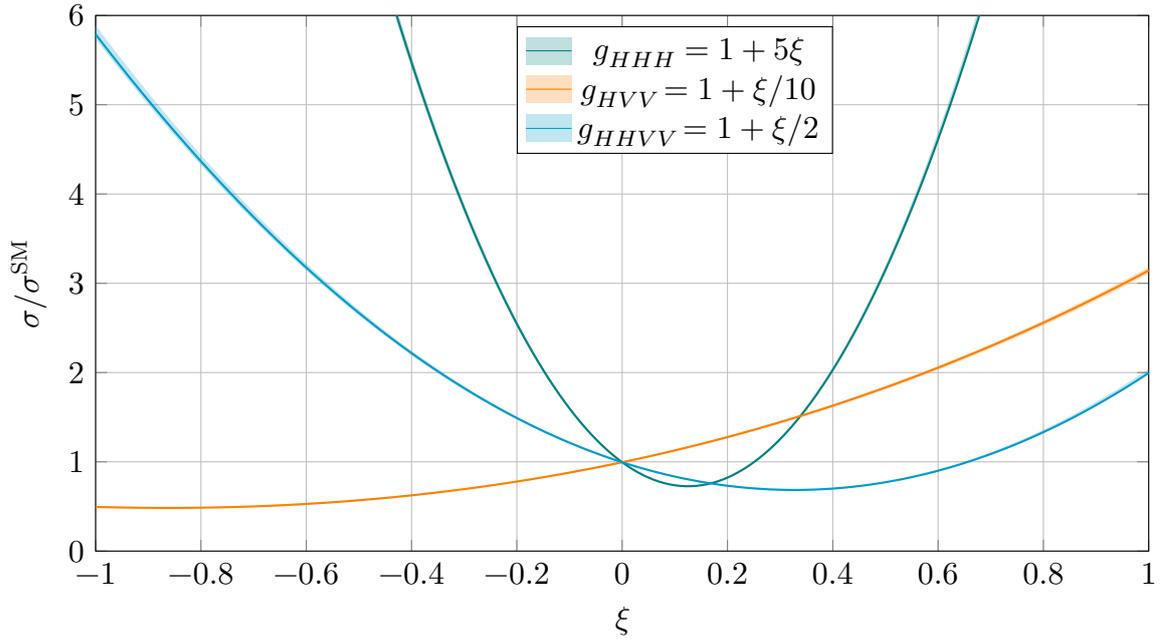
with coefficients  $\alpha$ ,  $\beta$  and  $\gamma$ . This form for the amplitude also holds at NLO QCD. The total cross section is then proportional to this expression squared, yielding the parameterization

$$\begin{aligned} \frac{\sigma}{\sigma_{\text{SM}}} = & A_0 g_{\text{HHH}}^2 g_{\text{HVV}}^2 + A_1 g_{\text{HVV}}^4 + A_2 g_{\text{HHV}}^2 \\ & + A_3 g_{\text{HHH}} g_{\text{HVV}}^3 + A_4 g_{\text{HHH}} g_{\text{HVV}} g_{\text{HHV}} + A_5 g_{\text{HVV}}^2 g_{\text{HHV}}. \end{aligned} \quad (4.14)$$

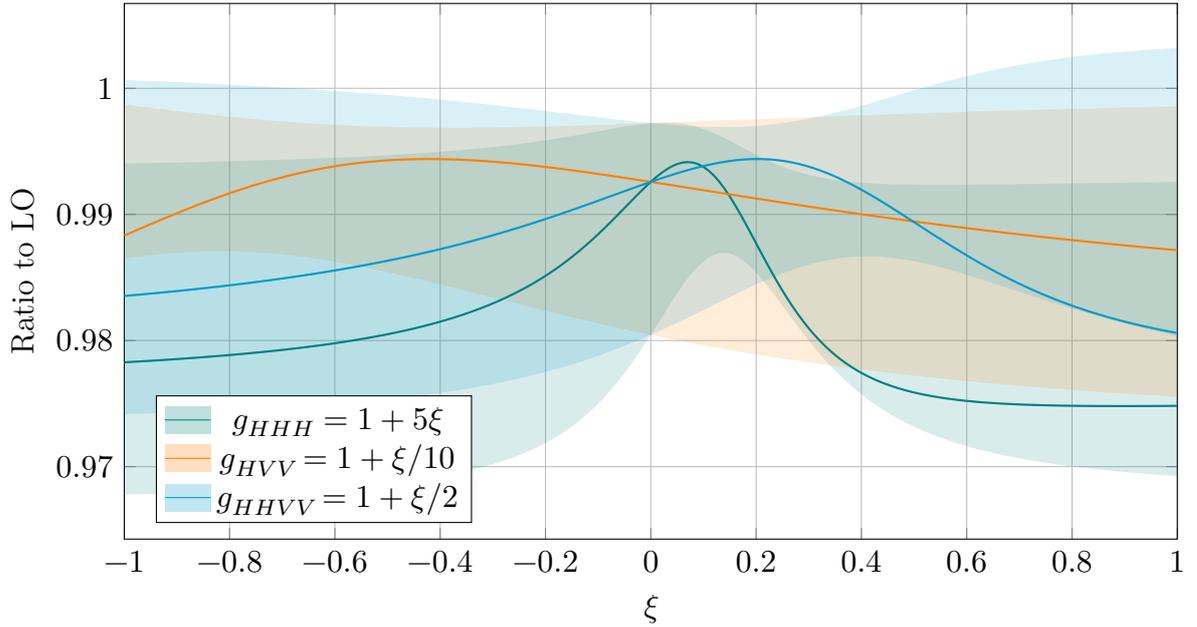
with coefficients  $A_i$ . This formula is exact, which is why the coefficients could in principle be determined by choosing six points of table 4.1 and analytically inverting the resulting linear system. However, due to the MC error on the cross sections, this results in large uncertainties on the coefficients. The simplest method to reduce this is by including all points of table 4.1, which is done with a fit. **This is performed with `iminuit` [56, 57] separately for each considered scale**, the resulting coefficients are shown in table 4.2.

This parameterization can now be used to investigate the dependence of the total cross sections on the anomalous couplings in the whole considered parameter space. Figure 4.5 shows the dependence of the cross section on a single anomalous coupling, where all other couplings are kept at the SM value. The impact of the different parameters varies strongly, with  **$g_{\text{HHH}}$  having by far the largest impact**, while  $g_{\text{HVV}}$  has the smallest impact. This is mostly a result of the size of the respective parameter ranges though. The vector boson couplings are far more constrained than the triple-Higgs coupling, allowing the latter to deviate much more from the SM value. Nonetheless though, all couplings can significantly alter the cross section.

### 4.3 Parameterization of the Total Cross Section in Terms of the Anomalous Couplings

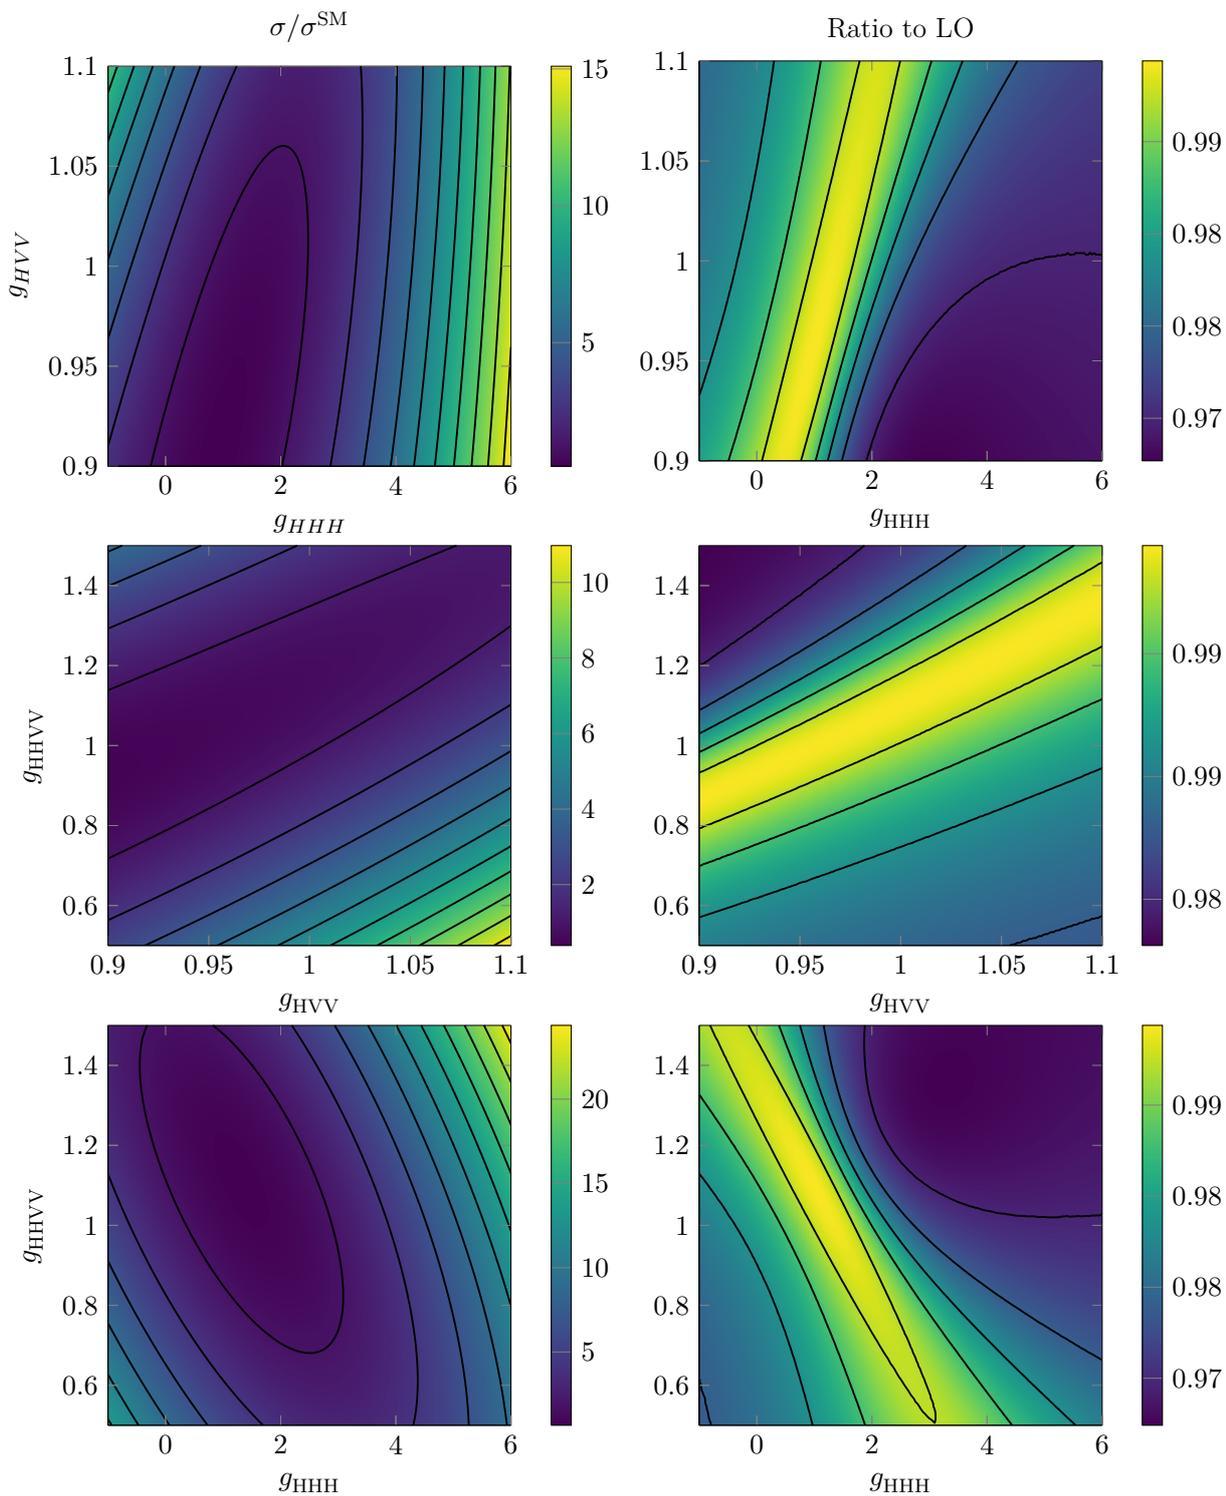


**Figure 4.5:** Total cross section at NLO relative to the Standard Model value when a single anomalous coupling is varied. The uncertainty band shows the scale uncertainty.

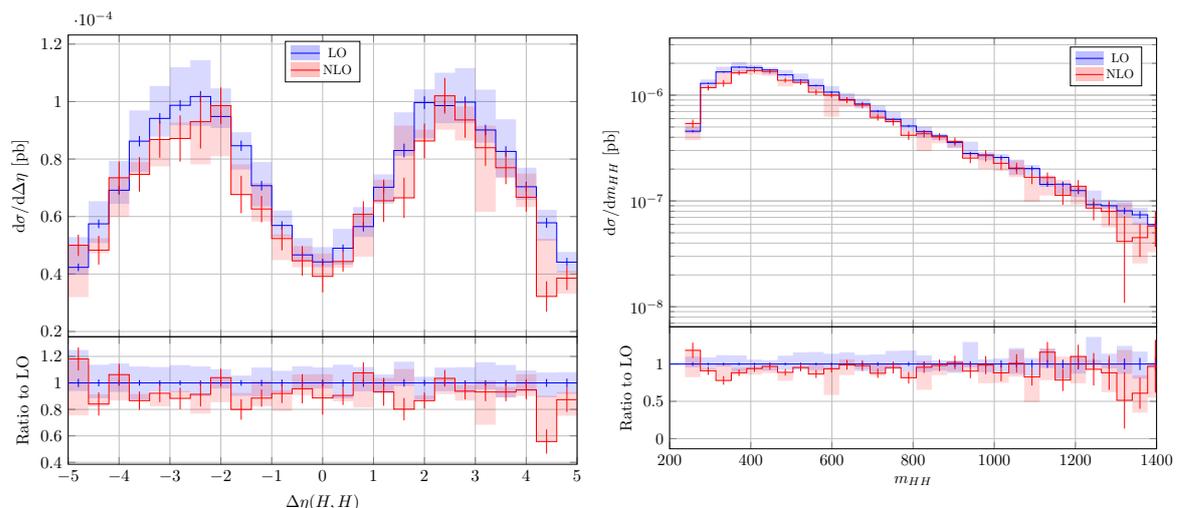


**Figure 4.6:** Ratio of the normalized cross section to the LO values,  $(\sigma_{\text{NLO}}/\sigma_{\text{NLO}}^{\text{SM}})/(\sigma_{\text{LO}}/\sigma_{\text{LO}}^{\text{SM}})$ . The uncertainty band shows the scale uncertainty.

Looking at the ratio to the LO values  $(\sigma_{\text{NLO}}/\sigma_{\text{NLO}}^{\text{SM}})/(\sigma_{\text{LO}}/\sigma_{\text{LO}}^{\text{SM}})$  in figure 4.6 next, the picture is as argued before: the  $K$ -factor is mostly flat. Varying a single parameter over its whole parameter space only impacts the ratio at most  $\sim 2\%$ , with the completely flat  $K$ -factor being included in the scale band for each coupling.



**Figure 4.7:** Total cross section relative to the SM at NLO (left) and ratio of this quantity to the LO value (right) in the  $g_{\text{HVV}}-g_{\text{HHH}}$  plane (top), the  $g_{\text{HHVV}}-g_{\text{HVV}}$  plane (center) and the  $g_{\text{HHVV}}-g_{\text{HHH}}$  plane (bottom).

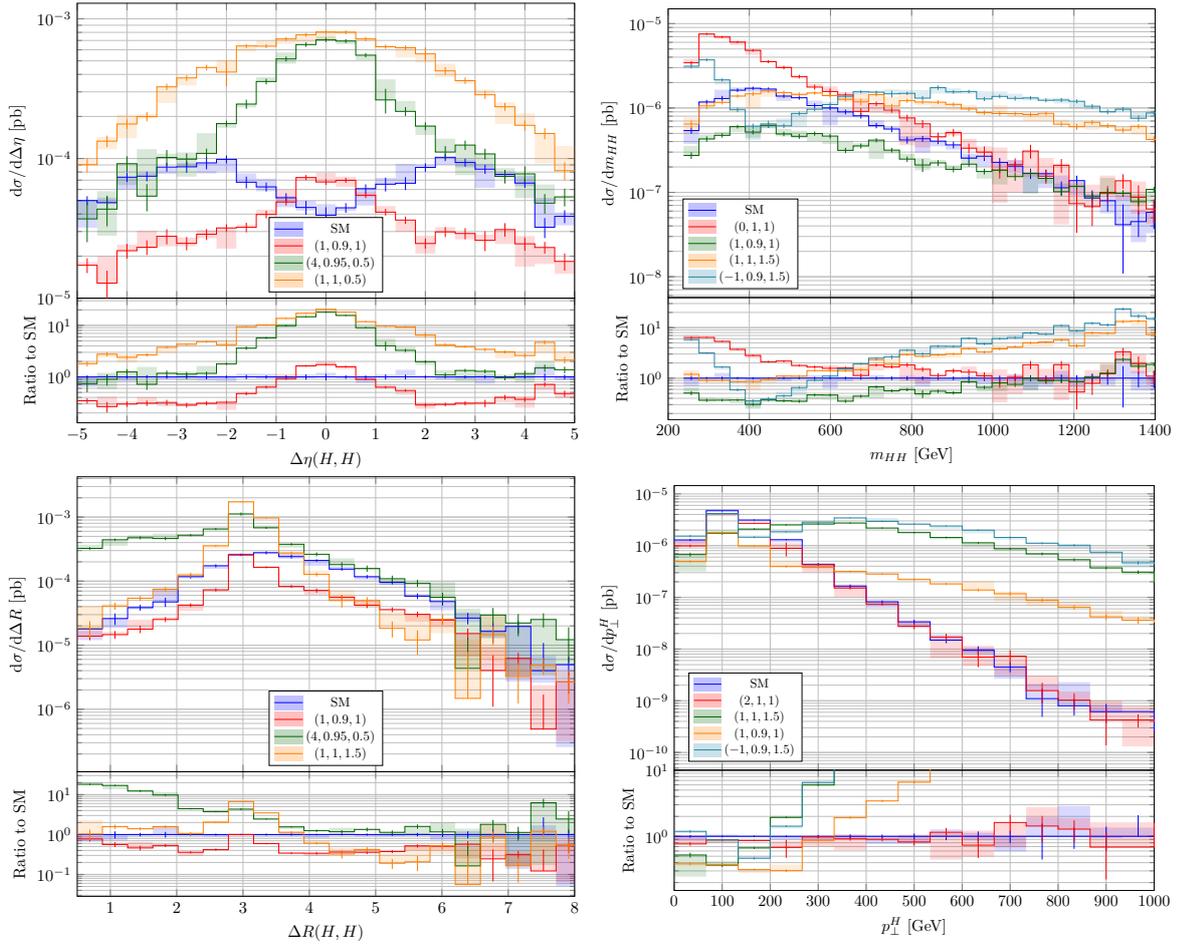


**Figure 4.8:** Pseudorapidity separation and invariant mass distributions of the di-Higgs system at LO and NLO in the SM. The error bars show the MC error, the error bands show the uncertainties due to the scale variation.

So far, only variations of a single coupling were considered. In figure 4.7, the simultaneous variation of two couplings is depicted. The left column contains the total cross section relative to the SM value at NLO, the right column the ratio of the cross section to the LO value. Each row then shows the respective quantity in the  $g_{\text{HVV}}-g_{\text{HHH}}$  plane, the  $g_{\text{HHVV}}-g_{\text{HVV}}$  plane and the  $g_{\text{HHVV}}-g_{\text{HHH}}$  plane. In each case, the third coupling is set to its SM value again. The  $K$ -factor is still mostly flat, in agreement with the previous observations. The largest deviations from the LO is for very large values of  $g_{\text{HHH}}$  and simultaneously small  $g_{\text{HVV}}$  or large  $g_{\text{HHVV}}$ , which is of order  $\sim 3\%$ . The strong dependence of the cross sections on the couplings can be seen in the left column, which shows deviations of more than 20 times the SM value for large  $g_{\text{HHH}}$  and  $g_{\text{HHVV}}$ . There are also possible deviations in the other direction though, will a roughly 50% decrease of the total cross section for small  $g_{\text{HVV}}$ .

## 4.4 Differential Results

Finally, several differential distributions are produced for each combination of anomalous couplings given in table 4.1. For this purpose, NLO events are generated at  $\sqrt{s} = 13.6$  TeV with WHIZARD and then analyzed with RIVET [58]. But before considering any anomalous couplings, some LO distributions can be compared to the respective NLO one. This is shown in the SM for the rapidity separation  $\Delta\eta = \eta(H_1) - \eta(H_2)$  of the two Higgs bosons and the di-Higgs invariant mass  $m_{\text{HH}} = \sqrt{(p_{H_1} + p_{H_2})^2}$  in figure 4.8. Both of these distributions show only mild shape differences between the LO and NLO histograms, which can be seen as representative for all generated distributions. The NLO QCD corrections mostly manifest themselves in the overall normalization due to the smaller cross section. They also show the expected decrease of the scale



**Figure 4.9:** Rapidity separation, invariant mass and  $R$ -separation distributions of the di-Higgs system and  $p_{\perp}$  distribution of any Higgs boson at NLO QCD. The anomalous couplings are given as  $(g_{HHH}, g_{HVV}, g_{HHVV})$  in the legend. The error bars show the MC error, the error bands show the uncertainties due to the scale variation.

uncertainty when going to NLO QCD, in accordance with the observation of the scale uncertainty on the total cross section in section 4.3.

Much larger shape deviations can be observed when the couplings are varied, as can be seen in figure 4.9. Here, distributions of several observables of the di-Higgs system are shown. **The strongest deviation from the SM distribution can be observed in the pseudorapidity separation of the Higgs bosons.** In the SM, the most Higgs boson pairs are produced with a pseudorapidity separation of  $\sim 2.5$  and a clear local minimum at  $\Delta\eta \sim 0$ . This is in contrast to the situation with anomalous couplings, where the distribution is skewed toward Higgs pairs with smaller pseudorapidity separation for variations in any of the three couplings. A clear local maximum forms in all shown scenarios for  $\Delta\eta \sim 0$ . This behavior can be understood by again considering the diagrams in figure 4.2. **The contribution of the leftmost diagram is enhanced for large self-couplings, leading to an excess in less separated Higgs pairs, since they originate in the same three-point vertex.** The situation is more subtle for the other

two couplings. As mentioned in e.g. [10], a cancellation happens between the second and third subdiagrams. If either of the two Higgs-vector boson couplings is varied, this cancellation fails and the **amplitude grows with energy**. This violates  $\mathcal{S}$ -matrix unitarity and has to be repaired by the UV-complete theory above the cutoff scale  $\Lambda$ . In the present observables however, this manifests in a large excess of high energy events, which can also clearly be seen in the  $p_{\perp}$  distribution of the Higgs bosons in figure 4.9. **This therefore also leads to the di-Higgs systems being more boosted on average, resulting in the smaller pseudorapidity separation.**

The changes in the pseudorapidity separation can also be observed in the  $R$ -separation  $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$  of the Higgs bosons, which is shown in the lower left histogram in figure 4.9. This distribution is, like the pseudorapidity separation, skewed towards smaller  $\Delta R$  values. Due to a simultaneous increase in  $\Delta\phi$  for anomalous values of the Higgs-vector boson couplings though, another interesting feature can be observed. **Increasing the self coupling only increases the abundance of low  $\Delta R$  events**, while a peak develops around  $\Delta R \sim 3$  for variations of the other two couplings. The large  $\Delta R$  region is essentially untouched by changes in the couplings.

Finally, figure 4.9 also contains the invariant mass  $m_{\text{HH}}$  of the di-Higgs system. The behavior with respect to the couplings is similar to before, changes in the Higgs-vector boson couplings manifest themselves in an excess of high energy events compared to the SM. The self-coupling on the other hand mostly influences low invariant mass events. By varying both at the same time, an interesting feature can be generated in the  $m_{\text{HH}}$  distribution: **a local minimum forms at  $m_{\text{HH}} \sim 450$  GeV with the coupling values shown in figure 4.9.**

In summary, variations of the three couplings  $g_{\text{HHH}}$ ,  $g_{\text{HVV}}$  and  $g_{\text{HHVV}}$  within their current experimental bounds can have a substantial influence on the cross sections and shapes of distributions in the VBF process. Additionally, completely new features can develop in the distributions, with some depending on the interplay of multiple anomalous couplings.



## 5 Summary and Outlook

In this thesis, the NLO QCD corrections to the VBF production of two Higgs bosons at the LHC within the HEFT were calculated. For this purpose, the relevant operators at the considered order of the HEFT were identified for this process. This model was used to calculate the total cross section and some differential distributions for several values of the anomalous couplings. The total cross sections at these points were then used to fit the coefficients of a parameterization of said cross section for arbitrary coupling values. Additionally, to make all these calculations more feasible, several improvements were implemented in the OLP GOSAM to reduce the time taken to construct the amplitude and evaluate it for a given phase space point.

However, even though many improvements were implemented in GOSAM, there is still much possibility for further improvement. Two areas of interest in this context are caching of intermediate results and further reduction of unnecessary recalculations in the kinematics routines. One important optimization implemented within the scope of this thesis was the reuse of results for helicities sharing the same kinematics. The kinematics of crossed helicities are also connected by permutations and reflections of momenta though. For this reason, much of the existing kinematical results could be reused for other helicities in a future version of GOSAM, further reducing unnecessarily spent time.

Turning to the VBF process, only the hard subprocess of an LHC collision is considered presently. The Higgs bosons and the highly energetic partons are not what is measured in the detector at the end, though. Therefore, a first natural extension of this calculation would be to include Higgs decays. This could be done by using the present setup to generate events with in general off-shell Higgs bosons, which can then in a further step be decayed by another program. A similar post-processing can be done for the highly energetic final state jets, which could be evolved down to a lower energy scale with a parton shower.

Apart from adding more components of a full LHC collision, the process itself could also be calculated to higher accuracy. Currently, only the NLO QCD corrections are included, even though the NLO electroweak corrections in the SM are only slightly smaller. Doing this consistently would require a large additional effort though. GOSAM is unable to perform the electroweak renormalization automatically, which would have to be done manually. Also, as mentioned in 2.2.2, the NLO HEFT operators have to be included with the NLO electroweak corrections. This would therefore also increase the number of variable parameters significantly, and make the renormalization even more complex at the same time.

However, already at the present stage, significant deviations of the SM predictions can be observed for anomalous couplings within the current experimental bounds, leading to very promising prospects for future experimental analyses at the (HL-)LHC.

# Bibliography

- [1] The ATLAS Collaboration. “Observation of a New Particle in the Search for the Standard Model Higgs Boson with the ATLAS Detector at the LHC”. In: *Phys. Lett. B* 716.1 (2012). arXiv: 1207.7214 [hep-ex].
- [2] The CMS Collaboration. “Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC”. In: *Phys. Lett. B* 716.1 (2012). arXiv: 1207.7235 [hep-ex].
- [3] S. Navas et al. “Review of Particle Physics”. In: *Phys. Rev. D* 110.3 (2024).
- [4] B. Di Micco et al. “Higgs Boson Potential at Colliders: Status and Perspectives”. In: *Rev. Phys.* 5 (2020). arXiv: 1910.00012 [hep-ph].
- [5] Gudrun Heinrich, Jannis Lang, and Ludovic Scyboz. “SMEFT Predictions for  $gg \rightarrow hh$  at Full NLO QCD and Truncation Uncertainties”. In: *J. High Energ. Phys.* 2022.8 (2022). arXiv: 2204.13045 [hep-ph].
- [6] Daniel de Florian et al. “Anomalous Couplings in Higgs-boson Pair Production at Approximate NNLO QCD”. In: *JHEP* 09.9 (2021). arXiv: 2106.14050 [hep-ph].
- [7] G. Heinrich et al. “A Non-Linear EFT Description of  $gg \rightarrow HH$  at NLO Interfaced to POWHEG”. In: *JHEP* 10.10 (2020). arXiv: 2006.16877 [hep-ph].
- [8] G. Buchalla et al. “Higgs Boson Pair Production in Non-Linear Effective Field Theory with Full  $m_t$ -Dependence at NLO QCD”. In: *JHEP* 09.9 (2018). arXiv: 1806.05162 [hep-ph].
- [9] Wolfgang Kilian et al. “Multi-Higgs Production and Unitarity in Vector-Boson Fusion at Future Hadron Colliders”. In: *Phys. Rev. D* 101.7 (2020). arXiv: 1808.05534 [hep-ph].
- [10] Fady Bishara, Roberto Contino, and Juan Rojo. “Higgs Pair Production in Vector-Boson Fusion at the LHC and Beyond”. In: *Eur. Phys. J. C* 77.7 (2017). arXiv: 1611.03860 [hep-ph].
- [11] Frédéric A. Dreyer et al. “Precise Predictions for Double-Higgs Production via Vector-Boson Fusion”. In: *Eur. Phys. J. C* 80.11 (2020). arXiv: 2005.13341 [hep-ph].
- [12] Frédéric A. Dreyer and Alexander Karlberg. “Vector-Boson Fusion Higgs Pair Production at N<sup>3</sup>LO”. In: *Phys. Rev. D* 98.11 (2018). arXiv: 1811.07906 [hep-ph].
- [13] Frédéric A. Dreyer and Alexander Karlberg. “Fully Differential Vector-Boson Fusion Higgs Pair Production at Next-to-Next-to-Leading Order”. In: *Phys. Rev. D* 99.7 (2019). arXiv: 1811.07918 [hep-ph].

- [14] Ling Liu-Sheng et al. “NNLO QCD Corrections to Higgs Pair Production via Vector Boson Fusion at Hadron Colliders”. In: *Phys. Rev. D* 89.7 (2014). arXiv: 1401.7754 [hep-ph].
- [15] Terrance Figy. “Next-to-Leading Order QCD Corrections to Light Higgs Pair Production via Vector Boson Fusion”. In: *Mod. Phys. Lett. A* 23.24 (2008). arXiv: 0806.2200 [hep-ph].
- [16] Michael Edward Peskin and Daniel V. Schroeder. *An Introduction to Quantum Field Theory*. Boca Raton: CRC Press, Taylor & Francis Group, 1995.
- [17] Ilaria Brivio and Michael Trott. “The Standard Model as an Effective Field Theory”. In: *Phys. Rept.* 793 (2019). arXiv: 1706.08945 [hep-ph].
- [18] W. Buchmüller and D. Wyler. “Effective Lagrangian Analysis of New Interactions and Flavour Conservation”. In: *Nucl. Phys. B* 268.3-4 (1986).
- [19] B. Grzadkowski et al. “Dimension-Six Terms in the Standard Model Lagrangian”. In: *JHEP* 10.10 (2010). arXiv: 1008.4884 [hep-ph].
- [20] Brian Henning et al. “2, 84, 30, 993, 560, 15456, 11962, 261485, ...: Higher Dimension Operators in the SM EFT”. In: *JHEP* 08 (2017). arXiv: 1512.03433 [hep-ph].
- [21] Gerhard Buchalla, Oscar Cata, and Claudius Krause. “Complete Electroweak Chiral Lagrangian with a Light Higgs at NLO”. In: *Nucl. Phys. B* 880 (2014). arXiv: 1307.5017 [hep-ph].
- [22] Gerhard Buchalla, Oscar Cata, and Claudius Krause. “On the Power Counting in Effective Field Theories”. In: *Phys. Lett. B* 731 (2014). arXiv: 1312.5624 [hep-ph].
- [23] G Peter Lepage. “A New Algorithm for Adaptive Multidimensional Integration”. In: *J. Comput. Phys.* 27.2 (1978).
- [24] Wolfgang Kilian, Thorsten Ohl, and Jurgen Reuter. “WHIZARD: Simulating Multi-Particle Processes at LHC and ILC”. In: *Eur. Phys. J. C* 71.9 (2011). arXiv: 0708.4233 [hep-ph].
- [25] Gavin Cullen et al. “Automated One-Loop Calculations with GoSam”. In: *Eur. Phys. J. C* 72.3 (2012). arXiv: 1111.2034 [hep-ph].
- [26] G. Cullen et al. “GoSam-2.0: A Tool for Automated One-Loop Calculations within the Standard Model and Beyond”. In: *Eur. Phys. J. C* 74.8 (2014). arXiv: 1404.7096 [hep-ph].
- [27] J. Alwall et al. “The Automated Computation of Tree-Level and next-to-Leading Order Differential Cross Sections, and Their Matching to Parton Shower Simulations”. In: *JHEP* 07.7 (2014). arXiv: 1405.0301 [hep-ph].
- [28] Federico Buccioni et al. “OpenLoops 2”. In: *Eur. Phys. J. C* 79.10 (2019). arXiv: 1907.13071 [hep-ph].
- [29] Ansgar Denner, Jean-Nicolas Lang, and Sandro Uccirati. “Recola2: REcursive Computation of One-Loop Amplitudes 2”. In: *Comput. Phys. Commun.* 224 (2018). arXiv: 1711.07388 [hep-ph].

- 
- [30] Ansgar Denner, Jean-Nicolas Lang, and Sandro Uccirati. “NLO Electroweak Corrections in Extended Higgs Sectors with RECOLA2”. In: *JHEP* 07.7 (2017). arXiv: 1705.06053 [hep-ph].
- [31] Julien Baglio et al. “Release Note – VBFNLO 3.0”. In: *arXiv:2405.06990 [hep-ph]* (2024). arXiv: 2405.06990 [hep-ph].
- [32] J. Baglio et al. “VBFNLO: A Parton Level Monte Carlo for Processes with Electroweak Bosons – Manual for Version 3.0”. In: *arXiv:1107.4038 [hep-ph]* (2024). arXiv: 1107.4038 [hep-ph].
- [33] Céline Degrande et al. “UFO - The Universal FeynRules Output”. In: *Comput. Phys. Commun.* 183.6 (2012). arXiv: 1108.2040 [hep-ph].
- [34] Luc Darmé et al. “UFO 2.0 – The Universal Feynman Output Format”. In: *Eur. Phys. J. C* 83.7 (2023). arXiv: 2304.09883 [hep-ph].
- [35] Adam Alloul et al. “FeynRules 2.0 - A Complete Toolbox for Tree-Level Phenomenology”. In: *Comput. Phys. Commun.* 185.8 (2014). arXiv: 1310.1921 [hep-ph].
- [36] P. Nogueira. “Automatic Feynman Graph Generation”. In: *J. Comput. Phys.* 105.2 (1993).
- [37] Ben Ruijl, Takahiro Ueda, and Jos Vermaseren. “FORM Version 4.2”. In: *arXiv:1707.06453 [hep-ph]* (2017). arXiv: 1707.06453 [hep-ph].
- [38] G. Cullen, M. Koch-Janusz, and T. Reiter. “Spinney: A Form Library for Helicity Spinors”. In: *Comput. Phys. Commun.* 182.11 (2011). arXiv: 1008.0803 [hep-ph].
- [39] J. Ph Guillet, G. Heinrich, and J. F. von Soden-Fraunhofen. “Tools for NLO Automation: Extension of the golem95C Integral Library”. In: *Comput. Phys. Commun.* 185.6 (2014). arXiv: 1312.3887 [hep-ph].
- [40] P. Mastrolia et al. “Scattering AMplitudes from Unitarity-based Reduction Algorithm at the Integrand-level”. In: *JHEP* 08.8 (2010). arXiv: 1006.0710 [hep-ph].
- [41] Tiziano Peraro. “Ninja: Automated Integrand Reduction via Laurent Expansion for One-Loop Amplitudes”. In: *Comput. Phys. Commun.* 185.10 (2014). arXiv: 1403.1229 [hep-ph].
- [42] Pierpaolo Mastrolia, Edoardo Mirabella, and Tiziano Peraro. “Integrand Reduction of One-Loop Scattering Amplitudes through Laurent Series Expansion”. In: *JHEP* 06.6 (2012). arXiv: 1203.0291 [hep-ph].
- [43] A. van Hameren. “OneLoop: For the Evaluation of One-Loop Scalar Functions”. In: *Comput. Phys. Commun.* 182.11 (2011). arXiv: 1007.4716 [hep-ph].
- [44] Stuart I. Feldman. “Make — a Program for Maintaining Computer Programs”. In: *Softw Pract Exp* 9.4 (1979).
- [45] S. Alioli et al. “Update of the Binoth Les Houches Accord for a Standard Interface between Monte Carlo Tools and One-Loop Programs”. In: *Comput. Phys. Commun.* 185.2 (2015). arXiv: 1308.3462 [hep-ph].

- [46] T. Binoth et al. “A Proposal for a Standard Interface between Monte Carlo Tools and One-Loop Programs”. In: *Comput. Phys. Commun.* 181.9 (2010). arXiv: 1001.1307 [hep-ph].
- [47] J. M. Campbell and R. K. Ellis. “An Update on Vector Boson Pair Production at Hadron Colliders”. In: *Phys. Rev. D* 60.11 (1999). arXiv: hep-ph/9905386.
- [48] Vladyslav Shtabovenko, Rolf Mertig, and Frederik Orellana. “FeynCalc 10: Do Multiloop Integrals Dream of Computer Codes?” In: *Comput. Phys. Commun.* 306 (2024). arXiv: 2312.14089 [hep-ph].
- [49] Andy Buckley et al. “LHAPDF6: Parton Density Access in the LHC Precision Era”. In: *Eur. Phys. J. C* 75.3 (2015). arXiv: 1412.7420 [hep-ph].
- [50] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. “The Anti-K<sub>t</sub> Jet Clustering Algorithm”. In: *JHEP* 04.04 (2008). arXiv: 0802.1189 [hep-ph].
- [51] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. “FastJet User Manual”. In: *Eur. Phys. J. C* 72.3 (2012). arXiv: 1111.6097 [hep-ph].
- [52] ATLAS Collaboration. “Combination of Searches for Higgs Boson Pair Production in  $pp$  Collisions at  $\sqrt{s} = 13$  TeV with the ATLAS Detector”. In: *Phys. Rev. Lett.* 133.10 (2024). arXiv: 2406.09971 [hep-ex].
- [53] C. M. S. Collaboration. “Constraints on the Higgs Boson Self-Coupling from the Combination of Single and Double Higgs Boson Production in Proton-Proton Collisions at  $\sqrt{s} = 13$  TeV”. In: *arXiv:2407.13554 [hep-ex]* (2024). arXiv: 2407.13554 [hep-ex].
- [54] J. Baglio et al. “The Measurement of the Higgs Self-Coupling at the LHC: Theoretical Status”. In: *JHEP* 04.4 (2013). arXiv: 1212.5581 [hep-ph].
- [55] T. Han, G. Valencia, and S. Willenbrock. “Structure-Function Approach to Vector-Boson Scattering in  $Pp$  Collisions”. In: *Phys. Rev. Lett.* 69.23 (1992). arXiv: hep-ph/9206246.
- [56] Hans Dembinski et al. *Scikit-Hep/Iminuit*. Version v2.30.2. Zenodo, 2024.
- [57] F. James and M. Roos. “Minuit - a System for Function Minimization and Analysis of the Parameter Errors and Correlations”. In: *Comput. Phys. Commun.* 10.6 (1975).
- [58] C. Bierlich et al. “Robust Independent Validation of Experiment and Theory: Rivet Version 3”. In: *SciPost Phys.* 8.2 (2020). arXiv: 1912.05451 [hep-ph].