AGH University of Krakow

Faculty of Physics and Applied Computer Science

MODULE 5

SYMMETRY BREAKING AND THE HIGGS SECTOR

Supporting Lecture Notes for *The Standard Model*

The Guideline

The aim of this lecture note is to make visible how spontaneous symmetry breaking and the Higgs sector resolve a central structural problem of the Standard Model: how massive weak gauge bosons and massive fermions can arise without abandoning gauge symmetry. Starting from the distinction between symmetry of the Lagrangian and symmetry of the vacuum, the note develops the logic of scalar symmetry breaking, Goldstone modes, the gauge-theory Higgs mechanism, the electroweak Higgs doublet, photon–Z mixing, and Yukawa mass generation. The later part of the note then turns to the physical Higgs boson, its couplings, and its collider realization at the LHC through selected representative Standard Model production channels and discovery signatures. Throughout the note, the emphasis is on seeing the Higgs sector both as a mathematically constrained part of the electroweak Standard Model and as the bridge between gauge structure and experimentally observed masses and Higgs phenomena.

Prepared for the Standard Model course

Hamzeh Khanpour

hamzeh.khanpour@fis.agh.edu.pl

https://view.fis.agh.edu.pl/staff/hamzeh.khanpour/

June 8, 2026

Contents

Preface

This supporting note is written for the *Module 5: Symmetry Breaking and the Higgs Sector* of the course *The Standard Model*. The guiding question of the module is simple to state but conceptually deep: how can the electroweak Standard Model describe the observed world of massive weak bosons and massive fermions if local gauge symmetry is one of its central organizing principles? The Higgs sector answers this question not by abandoning symmetry, but by realizing it in a subtler way. The Lagrangian is constructed to respect the electroweak gauge symmetry, while the vacuum state chosen by the theory does not share the full symmetry.

The pedagogical logic of the module can therefore be summarized as

$$\begin{aligned}&\text{electroweak gauge structure} \longrightarrow \text{mass problem} \longrightarrow \text{spontaneous symmetry breaking}\\ &\longrightarrow \text{Higgs mechanism} \longrightarrow \text{gauge-boson and fermion masses}\\ &\longrightarrow \text{physical Higgs boson} \longrightarrow \text{Higgs production at the LHC.}\end{aligned}$$

This lecture note is intentionally placed after the construction-oriented development of gauge theory and after the QCD module. It therefore does not rebuild the full Standard Model from the beginning. Instead, it focuses on one decisive structural issue. The same course logic that was used earlier, namely that symmetry, strains the allowed field content and interaction terms, is now pushed one step further: the mass pattern of the observed electroweak world must itself be understood as a consequence of a constrained quantum field theory.

The note has two closely related aims. The first is formal and conceptual: to show how spontaneous symmetry breaking works in scalar field theory, how the Goldstone theorem is modified when the symmetry is local, and how the electroweak Higgs doublet gives masses to the $W^\pm$, Z, and the fermions while leaving the photon massless. The second aim is phenomenological but still bounded: to connect this formal structure to the observed Higgs boson at the LHC and to the main Standard Model production and decay signatures at the LHC. This later part is included not to turn Module 5 into a full collider-physics chapter, but to make the theory visibly continuous with experiment. The collider-facing discussion is therefore intentionally qualitative, selective, and bridge-like, pointing forward to later phenomenology modules without trying to replace them.

Conventions and notation

- Metric signature: $\eta_{\mu\nu} = \mathrm{diag}(1, -1, -1, -1)$.
- Natural units: $\hbar = c = 1$.
- Electroweak gauge group before symmetry breaking: $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$.
- Electric charge operator: $Q = T_3 + Y$.
- Higgs doublet notation:

$$\Phi = \begin{pmatrix}\phi^+\\ \phi^0\end{pmatrix}, \qquad Y_\Phi = +\tfrac{1}{2}.$$

- Electroweak covariant derivative:

$$D_\mu = \partial_\mu - \mathrm{i}g\,\frac{\tau^a}{2}W^a_\mu - \mathrm{i}g' \,Y B_\mu.$$

- Higgs potential:

$$V(\Phi) = \mu^2\Phi^\dagger\Phi + \lambda(\Phi^\dagger\Phi)^2, \qquad \lambda > 0.$$

- Vacuum expectation value (vev) notation:

$$\langle\Phi\rangle = \frac{1}{\sqrt{2}}\begin{pmatrix}0\\ v\end{pmatrix}.$$

- Weak mixing angle notation: θ_W.
- Charged gauge-boson combinations:

$$W^\pm_\mu = \frac{1}{\sqrt{2}}\left(W^1_\mu \mp \mathrm{i}W^2_\mu\right).$$

1 Introduction and module roadmap

1.1 Why the Higgs sector deserves its own module

The Higgs sector deserves its own module because the question of mass generation is not a technical afterthought added late in the Standard Model. It is one of the places where the internal logic of the theory is most sharply tested. In a non-gauge theory one may write mass terms almost immediately, provided that Lorentz invariance is respected. In the electroweak Standard Model the situation is different. Local gauge invariance and chiral matter assignments are so restrictive that the masses of weak gauge bosons and fermions cannot simply be inserted by hand without damaging the very structure that made the theory predictive.

This means that Module 5 is not merely "the Higgs boson chapter". Its real purpose is to show how the Standard Model can remain a gauge theory while still matching the observed world of massive $W^\pm$, Z, quarks, and charged leptons. Historically, the Higgs boson entered the subject through this structural problem. Experiment later observed a scalar particle compatible with this mechanism, but the conceptual heart of the topic remains the same: the Higgs sector is a mechanism for reconciling symmetry with mass.

1.2 From Module 3 and Module 4 to Module 5

Module 3 established the general logic of gauge symmetry. It showed how local symmetry leads naturally to gauge fields, covariant derivatives, and tightly constrained interactions. The Standard Model was then organized as a gauge theory based on $\mathrm{SU}(3)_c \times \mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$, with matter fields arranged into chiral multiplets and with their interaction content fixed by gauge invariance.

Module 4 then made one important part of this logic concrete by following the $\mathrm{SU}(3)_c$ sector into QCD. There, the abstract structure of a non-Abelian gauge theory became visible in physical consequences such as gluon self-interactions, asymptotic freedom, confinement, hadrons, jets, and collider phenomena. Module 5 now plays the analogous role for the electroweak sector. The guiding question changes: not "how does a gauge theory describe the strong interaction?", but rather "how does an electroweak gauge theory become the observed massive weak sector?" Photon–Z mixing, the masses of the $W^{\pm}$ and Z, and the Yukawa origin of fermion masses are therefore natural continuations of the construction logic already established.

1.3 What this module will and will not do

This module will develop spontaneous symmetry breaking in scalar field theory, explain the distinction between symmetry of the Lagrangian and symmetry of the vacuum, introduce the Higgs mechanism first in prototype form and then in the electroweak Standard Model, derive the masses of the weak gauge bosons, explain photon–Z mixing, construct Yukawa mass generation for fermions, identify the physical Higgs boson, and connect the formalism to the observed Higgs boson at the LHC through a bounded first look at representative Standard Model production and discovery channels.

At the same time, the module has deliberately bounded scope. It will not provide a full treatment of general gauge fixing in R_ξ gauges, a full precision-Higgs-coupling programme, a complete flavour analysis with CKM diagonalisation, a detailed EFT treatment of extended scalar sectors, or a full experimental collider-analysis manual. Those topics either belong to later modules or to more advanced courses.

1.4 Roadmap of the note

The note begins by formulating the electroweak mass problem in a sharp way. It then introduces spontaneous symmetry breaking in simple scalar theories, first for a discrete symmetry and then for a continuous global symmetry, so that the reader can see clearly what changes when the symmetry becomes local. After the Abelian Higgs model has illustrated the Higgs mechanism in prototype form, the note returns to the Standard Model and develops the electroweak Higgs doublet, the scalar potential, the vacuum expectation value, degree-of-freedom counting, gauge-boson masses, photon–Z mixing, and Yukawa interactions. The later sections then turn to the physical Higgs boson itself, its couplings and basic observable properties, and finally to a bounded bridge to its collider realization at the LHC.

Compact roadmap. The logical flow of the module is

$$\begin{aligned}&\text{mass problem} \longrightarrow \text{spontaneous symmetry breaking} \longrightarrow \text{Higgs mechanism}\\&\longrightarrow \text{electroweak symmetry breaking} \longrightarrow \text{masses and couplings}\\&\longrightarrow \text{Higgs boson} \longrightarrow \text{bounded bridge to the LHC.}\end{aligned}$$

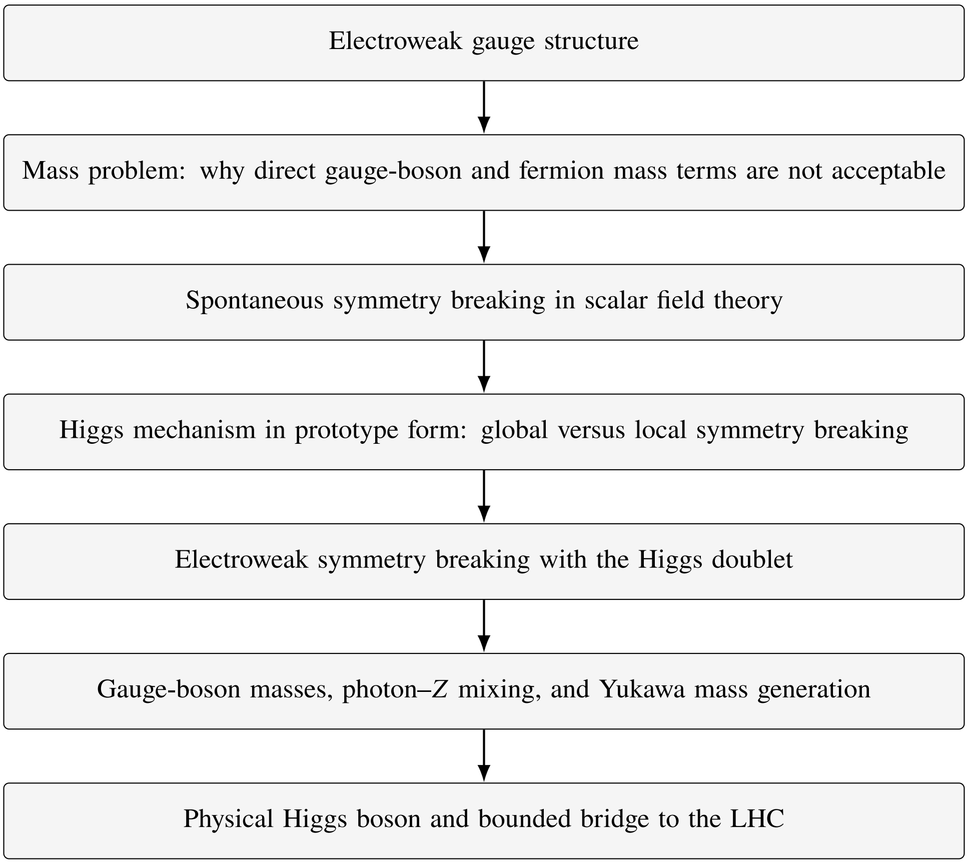

Figure 1: Conceptual roadmap of Module 5. The note begins with the electroweak mass problem, develops spontaneous symmetry breaking and the Higgs mechanism in prototype form, and then applies this logic to electroweak symmetry breaking, mass generation, the physical Higgs boson, and its bounded experimental interpretation.

2 Why masses are nontrivial in the electroweak theory

2.1 The empirical fact of massive weak bosons and massive fermions

Experiment tells us immediately that the weak interaction is not mediated by massless gauge bosons. The observed weak bosons are the charged $W^{\pm}$ and the neutral Z, with masses of order 10^2 GeV, whereas the photon is massless. The observed fermions are also massive, with masses spanning many orders of magnitude from the electron to the top quark. At a purely empirical level these are elementary facts. At the level of theory, however, they pose a structural question: where do these masses come from inside a gauge theory whose field content and interaction terms were fixed before any masses were written down?

This question is sharpened by the chiral nature of the electroweak theory. Left-handed and right-handed fermions do not transform in the same way under $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$. The weak interaction distinguishes chirality. Therefore the mass problem is not merely a question of introducing a scalar field with a nonzero expectation value. It is specifically a question about how a chiral gauge theory can account for observed masses while preserving the gauge structure that makes the Standard Model coherent.

2.2 Why a direct gauge-boson mass term is not acceptable

Consider first an Abelian gauge field A_μ. A direct mass term would have Proca form,

$$\mathcal{L}_{m_A} = \tfrac{1}{2} m_A^2 A_\mu A^\mu. \tag{2.1}$$

In an Abelian gauge theory the gauge transformation is

$$A_\mu(x) \longrightarrow A_\mu(x) + \partial_\mu \alpha(x), \tag{2.2}$$

so the quantity $A_\mu A^\mu$ does not remain invariant. Under the transformation above it picks up terms proportional to $A^\mu \partial_\mu \alpha$ and $(\partial_\mu \alpha)(\partial^\mu \alpha)$. Thus a direct mass term destroys local gauge invariance.

The same issue persists, in more elaborate form, for non-Abelian gauge fields. A direct mass term schematically of the form $m^2 W_\mu^a W^{a\mu}/2$ is not invariant under the local $SU(2)_L$ gauge symmetry. This is not a minor aesthetic complaint. Gauge symmetry was the principle that determined the theory's covariant derivative, field strength, and interaction structure. If one simply inserts masses in a way that violates it, the theory ceases to be the electroweak gauge theory one intended to write.

2.3 Why a direct fermion mass term is not acceptable in a chiral gauge theory

For a Dirac fermion ψ, a naive mass term would be

$$\mathcal{L}_{m_f} = -m\bar{\psi}\psi. \tag{2.3}$$

Decomposing ψ into left- and right-handed parts,

$$\psi = \psi_L + \psi_R, \qquad \psi_{L,R} = \frac{1 \mp \gamma^5}{2}\psi, \tag{2.4}$$

we may write

$$\bar{\psi}\psi = \bar{\psi}_L \psi_R + \bar{\psi}_R \psi_L. \tag{2.5}$$

Equation (2.5) makes the electroweak difficulty visible. In the Standard Model, ψ_L and ψ_R generally belong to different representations of $SU(2)_L \times U(1)_Y$. For example, the charged-lepton left-handed field sits in a doublet

$$L_L = \begin{pmatrix} \nu_L \\ e_L \end{pmatrix}, \tag{2.6}$$

whereas e_R is an $SU(2)_L$ singlet. Since the two pieces transform differently, the product $\bar{\psi}_L \psi_R$ is not, in general, a gauge singlet. Therefore a naive fermion mass term is forbidden before electroweak symmetry breaking.

This is one of the deepest reasons why the Higgs sector is needed. It is not merely added to endow bosons with mass; it is also required to construct gauge-invariant couplings that connect left- and right-handed fermion fields.

Example 2.1: Why the naive electron mass term fails

Before electroweak symmetry breaking, the field e_L is not an isolated electroweak singlet: it is the lower component of the weak doublet

$$L_L = \begin{pmatrix} \nu_e \\ e \end{pmatrix}_L ,$$

whereas e_R is an $SU(2)_L$ singlet. A term such as $\bar{e}_L e_R$ therefore does not, by itself, transform as an $SU(2)_L$ singlet. This is the group-theoretic reason why a direct electron mass term is not allowed before the Higgs field is introduced.

2.4 A short consistency remark: gauge symmetry, renormalisability, and unitarity

At the level of a first course, it is enough to keep one bounded physical point in mind. Gauge symmetry is not merely a convenient language for organising interactions. It is deeply related to the high-energy consistency of the theory. If massive vector bosons were introduced by hand, scattering amplitudes involving their longitudinal polarization states would grow too strongly with energy. The Higgs sector modifies these amplitudes in precisely the way needed to tame this growth. In this sense the Higgs field is tied not only to mass generation but also to the perturbative consistency of the electroweak theory.

One should therefore read the Higgs mechanism as solving several problems at once. It preserves gauge invariance at the level of the Lagrangian, explains the observed masses of the weak bosons, removes physical massless Goldstone bosons from the spectrum in the gauged case, and contributes to the good high-energy behaviour of the theory.

2.5 Statement of the strategy: keep the Lagrangian symmetric, let the vacuum break the symmetry

The central strategy of the Higgs mechanism can now be stated succinctly. One does not write explicit mass terms that violate the gauge symmetry. Instead, one adds a scalar sector whose potential is such that the vacuum chooses a nonzero expectation value. The full Lagrangian remains gauge invariant, but the vacuum no longer exhibits the full symmetry. When the theory is expanded about this vacuum, the excitations reorganise themselves: some gauge bosons become massive, fermion masses can arise from Yukawa couplings, and one physical scalar field remains. The next sections build this logic step by step.

3 Symmetry breaking in scalar field theory: first examples

3.1 Symmetry of the Lagrangian versus symmetry of the vacuum

The distinction between symmetry of the Lagrangian and symmetry of the vacuum is the conceptual foundation of the whole module. A symmetry of the Lagrangian means that the equations governing the theory are invariant under a certain transformation. A symmetry of the vacuum means that the chosen ground state is itself invariant under that transformation. These two statements are not equivalent.

Definition 3.1: Symmetry of the vacuum

A vacuum is said to be symmetric under a transformation if the vacuum state is left invariant by that transformation. A theory may have a Lagrangian that is symmetric even when the chosen vacuum is not. Spontaneous symmetry breaking refers precisely to this mismatch between the symmetry of the equations and the symmetry of the ground state.

A theory may possess a symmetric Lagrangian but admit several degenerate minima. If the system settles into one specific minimum, then the vacuum need not share the symmetry of the Lagrangian. The symmetry has not been explicitly destroyed; it is still present in the equations. Rather, it is the choice of ground state that has singled out one direction in field space. This is what is meant by spontaneous symmetry breaking.

Remark 3.1: Spontaneous versus explicit symmetry breaking

In explicit symmetry breaking, the Lagrangian itself contains terms that fail to respect the symmetry. In spontaneous symmetry breaking, the Lagrangian remains symmetric, but the chosen vacuum does not display the full symmetry. This distinction is essential: the Higgs mechanism relies on the second case, not the first.

3.2 Spontaneous breaking of a discrete symmetry

A useful warm-up model is a single real scalar field ϕ with Lagrangian

$$\mathcal{L} = \tfrac{1}{2}\partial_\mu\phi\,\partial^\mu\phi - V(\phi), \tag{3.1}$$

where the potential is

$$V(\phi) = \frac{\lambda}{4}(\phi^2 - v^2)^2, \qquad \lambda > 0. \tag{3.2}$$

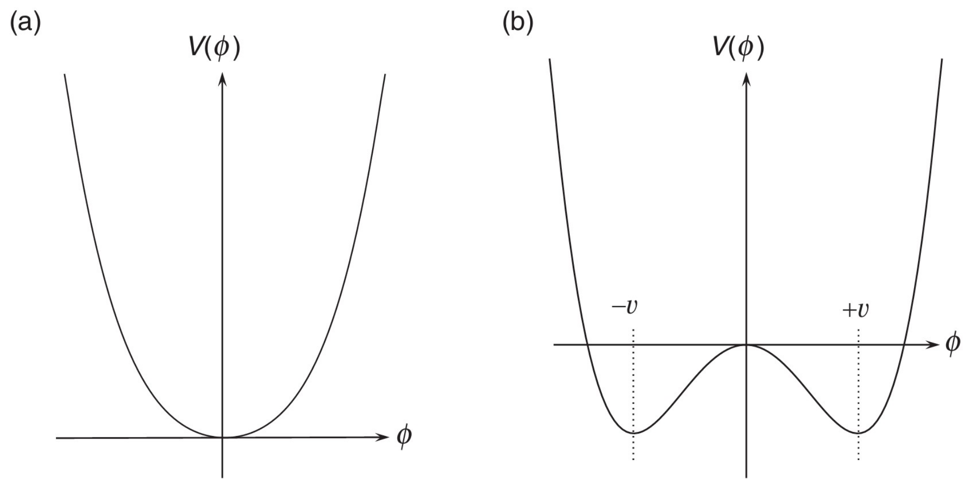

Figure 2: One-dimensional scalar potentials illustrating the transition from an unbroken phase with a unique minimum at $\phi = 0$ to a spontaneously broken phase with two degenerate minima at $\phi = \pm v$. This is the simplest field-theoretic example of how the symmetry of the Lagrangian can remain intact while the chosen vacuum does not share that symmetry.

This potential is invariant under the discrete symmetry $\phi \to -\phi$. The two minima are at

$$\phi = +v, \qquad \phi = -v. \tag{3.3}$$

The Lagrangian is symmetric, but once one chooses one of the two minima as the vacuum, that chosen vacuum is not invariant under $\phi \to -\phi$. The transformation maps one vacuum into the other.

This simple example already captures the essence of spontaneous symmetry breaking. Nothing was added to the Lagrangian to break the symmetry explicitly. Instead, the potential had more than one minimum, and the system had to choose one.

3.3 Vacuum degeneracy and choice of ground state

The existence of several symmetry-related minima is called vacuum degeneracy. For the discrete example above, there are two degenerate ground states. More generally, the vacuum manifold may consist of several isolated points or of a continuous set of minima. The physical theory around a chosen vacuum is obtained by expanding the fields about that vacuum. This means that practical calculations are performed not around the symmetric point $\phi = 0$, but around the selected ground state.

At this stage it is important to stress that choosing one vacuum does not mean the original symmetry has ceased to exist as a property of the Lagrangian. Rather, it means that the symmetry is not manifest in the particular ground state used to define the particle excitations. This difference between the laws and the ground state is what later allows the electroweak theory to preserve gauge symmetry in the Lagrangian while describing a massive low-energy world.

3.4 Expansion about the chosen vacuum

Choose the vacuum $\phi = v$. Write the field as

$$\phi(x) = v + \eta(x), \tag{3.4}$$

where η describes fluctuations about the chosen minimum. Substituting into (3.2), one finds

$$V(v + \eta) = \lambda v^2 \eta^2 + \lambda v \eta^3 + \frac{\lambda}{4}\eta^4. \tag{3.5}$$

Using the standard quadratic normalization $V \supset \frac{1}{2}m^2\eta^2$, the fluctuation field has mass

$$m_\eta^2 = 2\lambda v^2. \tag{3.6}$$

Example 3.1: Mass from expansion about the vacuum

Starting from $\phi(x) = v + \eta(x)$, the potential becomes

$$V(v+\eta) = \lambda v^2\eta^2 + \lambda v\eta^3 + \frac{\lambda}{4}\eta^4.$$

Comparing the quadratic term with the standard form $\frac{1}{2}m_\eta^2\eta^2$, one reads off

$$m_\eta^2 = 2\lambda v^2.$$

The important point is conceptual: the mass of the fluctuation field is not inserted by hand in the original symmetric form of the theory, but appears after expanding around the chosen vacuum.

Thus a massive scalar excitation appears after expanding about the vacuum. The important lesson is that the mass arises not by being written directly in the original symmetric form of the theory, but by expanding around a nontrivial vacuum.

4 Continuous global symmetry breaking and Goldstone modes

4.1 A complex scalar or $O(N)$-type prototype

To prepare for the electroweak case, we now consider a theory with a continuous global symmetry. The simplest example is a complex scalar field ϕ with Lagrangian

$$\mathcal{L} = \partial_\mu\phi^*\,\partial^\mu\phi - \frac{\lambda}{2}\left(|\phi|^2 - \frac{v^2}{2}\right)^2. \tag{4.1}$$

This theory is invariant under the global U(1) transformation

$$\phi(x) \longrightarrow \mathrm{e}^{\mathrm{i}\alpha}\phi(x), \tag{4.2}$$

with constant α. The minima satisfy

$$|\phi|^2 = \frac{v^2}{2}, \tag{4.3}$$

so the set of minima forms a circle in the two-dimensional field space of $\Re\phi$ and $\Im\phi$.

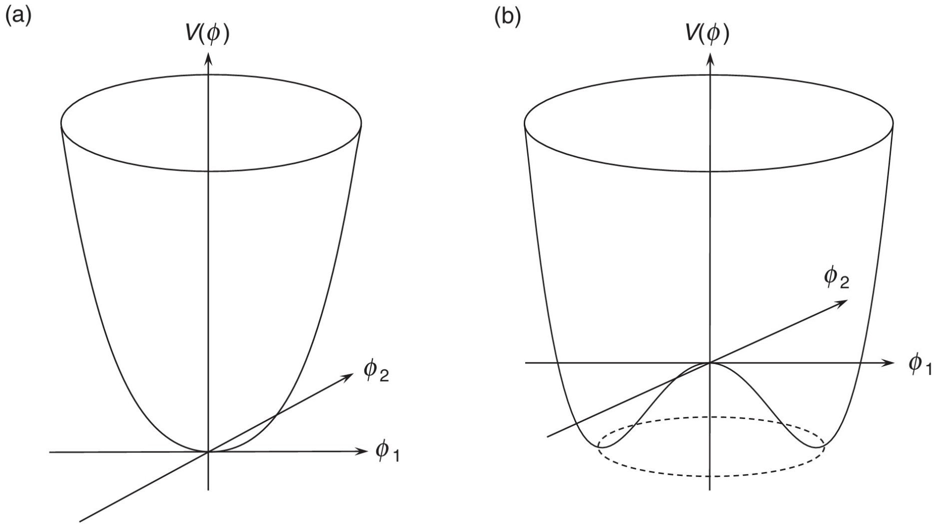

Figure 3: Potential for a complex scalar field with global $U(1)$ symmetry. In the broken case, the minima form a continuous circle in field space rather than isolated points, making visible the vacuum manifold on which the Goldstone mode lives.

An equivalent pedagogical language is the $O(N)$ model, in which a multiplet of real scalars is rotated among itself by a continuous global symmetry. For the present module the complex scalar example already captures the essential physics while keeping the algebra transparent.

4.2 Broken generators and the vacuum manifold

Unlike the discrete example, the vacuum manifold is now continuous. Every point on the circle $|\phi| = v/\sqrt{2}$ is a classical ground state, and the global $\mathrm{U}(1)$ transformation moves one vacuum into another. Choosing a specific vacuum therefore breaks the symmetry in the vacuum, even though the Lagrangian remains invariant.

A convenient parametrisation is

$$\phi(x) = \frac{1}{\sqrt{2}}\,(v + h(x))\,\mathrm{e}^{\mathrm{i}\pi(x)/v}. \tag{4.4}$$

The field $h(x)$ points radially away from the vacuum manifold, while $\pi(x)$ moves tangentially around it. This decomposition makes the geometry of the vacuum manifold visible: radial fluctuations change $|\phi|$, whereas angular fluctuations move along the degenerate set of vacua.

4.3 Goldstone's theorem at course level

The basic statement of Goldstone's theorem is the following: if a continuous global symmetry is spontaneously broken, then the spectrum contains one massless scalar mode for each broken generator. In the present $\mathrm{U}(1)$ example there is one broken generator, so one expects one massless mode.

Definition 4.1: Goldstone mode

A Goldstone mode is a massless scalar excitation associated with the spontaneous breaking of a continuous global symmetry. Geometrically, it corresponds to motion along the vacuum manifold, that is, along directions in field space that do not change the potential energy at quadratic order.

For the purposes of this course, the theorem will be used mainly as a physical result rather than proved in full quantum-field-theoretic detail. The important point is that the existence of the massless mode is not accidental. It follows from the fact that motion along the vacuum manifold costs no potential energy at quadratic order.

4.4 Why massless Goldstone modes appear in the global case

Inserting the parametrisation (4.4) into (4.1), one finds

$$\mathcal{L} = \tfrac{1}{2}(\partial_\mu h)(\partial^\mu h) + \tfrac{1}{2}\left(1 + \frac{h}{v}\right)^2 (\partial_\mu \pi)(\partial^\mu \pi) - V(h), \tag{4.5}$$

where

$$V(h) = \frac{\lambda}{8}\left((v + h)^2 - v^2\right)^2 . \tag{4.6}$$

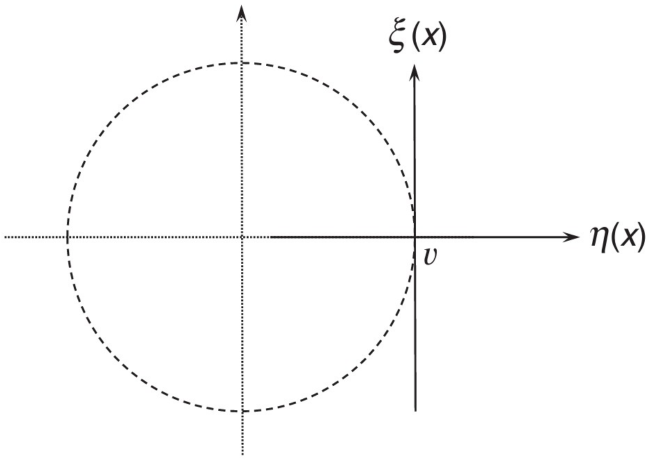

Figure 4: Field decomposition around a chosen vacuum in the broken $U(1)$ theory. The radial fluctuation $\eta(x)$ corresponds to a massive scalar mode, while the angular fluctuation $\xi(x)$ moves along the vacuum manifold and corresponds to the massless Goldstone mode.

At quadratic order, the field h is massive while π appears only through derivative terms and has no mass term. Thus the radial mode is massive and the angular mode is massless. The massless field π is the Goldstone boson of the broken global symmetry.

The global case therefore has a clear lesson: spontaneous breaking of a continuous global symmetry produces massless scalar excitations. The electroweak theory cannot end at this stage, because no such massless Goldstone bosons are observed as physical particles associated with electroweak symmetry breaking.

Guided checks

- In the parametrisation $\phi(x) = \frac{1}{\sqrt{2}}(v + h(x))e^{i\pi(x)/v}$, which field corresponds to radial motion and which corresponds to angular motion?
- Why does motion along the vacuum manifold fail to generate a mass term at quadratic order?
- How many Goldstone modes would you expect if two independent continuous generators were broken?
- Why is the global $U(1)$ example not yet sufficient to describe electroweak symmetry breaking?

4.5 Why this is only a stepping-stone to the electroweak story

The Goldstone theorem is not the end of the symmetry-breaking story, but the beginning of the crucial distinction between global and local symmetry. The electroweak symmetry is gauged, not merely global. In a gauge theory, field configurations related by gauge transformations describe the same physics. This changes the interpretation of the vacuum manifold and the status of the would-be Goldstone modes. In particular, the same angular directions that gave physical massless modes in the global case will, in the gauged case, become tied to gauge redundancy and to the longitudinal degrees of freedom of massive vector bosons. The next section explains this change in the simplest possible setting.

5 From global to local symmetry: the Higgs mechanism in prototype form

5.1 Why the gauged case changes the conclusion

When the symmetry is local, the transformation parameter becomes a function of spacetime. For a U(1) gauge theory one has

$$\phi(x) \longrightarrow \mathrm{e}^{\mathrm{i}e\alpha(x)}\phi(x), \qquad A_\mu(x) \longrightarrow A_\mu(x) + \partial_\mu\alpha(x). \tag{5.1}$$

Field configurations connected by such a gauge transformation are physically equivalent. In particular, motion around the vacuum manifold can be reinterpreted as gauge redundancy rather than as a physically distinct massless excitation. The key results are then the opposite of the global case: there is no physical massless Goldstone boson, while the gauge field becomes massive.

5.2 The Abelian Higgs model

The simplest realisation is the Abelian Higgs model,

$$\mathcal{L} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + (D_\mu\phi)^*(D^\mu\phi) - \frac{\lambda}{2}\left(|\phi|^2 - \frac{v^2}{2}\right)^2, \tag{5.2}$$

with

$$D_\mu\phi = (\partial_\mu - \mathrm{i}eA_\mu)\phi, \qquad F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu. \tag{5.3}$$

The vacuum again satisfies $|\phi| = v/\sqrt{2}$. Parametrise the field as

$$\phi(x) = \frac{1}{\sqrt{2}}(v + h(x))\mathrm{e}^{\mathrm{i}\xi(x)/v}. \tag{5.4}$$

Then

$$D_\mu \phi = \frac{\mathrm{e}^{\mathrm{i}\xi/v}}{\sqrt{2}}\left[\partial_\mu h + \mathrm{i}(v + h)\left(\frac{1}{v}\partial_\mu \xi - eA_\mu\right)\right]. \tag{5.5}$$

The combination $(\partial_\mu \xi/v - eA_\mu)$ shows that the would-be Goldstone field ξ is tied directly to the gauge field.

5.3 Unitary gauge and the disappearance of the Goldstone mode

A convenient choice is unitary gauge, obtained by using the local gauge freedom to set $\xi(x) = 0$. The scalar field then becomes simply

$$\phi(x) = \frac{1}{\sqrt{2}}(v + h(x)). \tag{5.6}$$

Substituting into (5.2) yields, among other terms,

$$\mathcal{L} \supset -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + \tfrac{1}{2}e^2v^2A_\mu A^\mu + \tfrac{1}{2}(\partial_\mu h)(\partial^\mu h) - \tfrac{1}{2}(2\lambda v^2)h^2 + \cdots. \tag{5.7}$$

The field A_μ has acquired a mass

$$m_A = ev, \tag{5.8}$$

while the scalar field h remains as a massive physical excitation. The Goldstone mode ξ has disappeared from the physical spectrum. It has not become unimportant; its degree of freedom has been absorbed into the longitudinal polarization of the now massive gauge field.

5.4 Massive vector boson plus one remaining scalar

Degree-of-freedom counting makes the mechanism transparent. Before symmetry breaking, the theory contains a massless gauge field with two physical transverse polarizations and a complex scalar with two real degrees of freedom. This totals four real degrees of freedom. After symmetry breaking in unitary gauge, the spectrum contains a massive vector boson with three polarizations and one real massive scalar, the Higgs field h. The total remains four. The Higgs mechanism is therefore not creation of extra degrees of freedom; it is a reorganisation of the spectrum around a nontrivial vacuum.

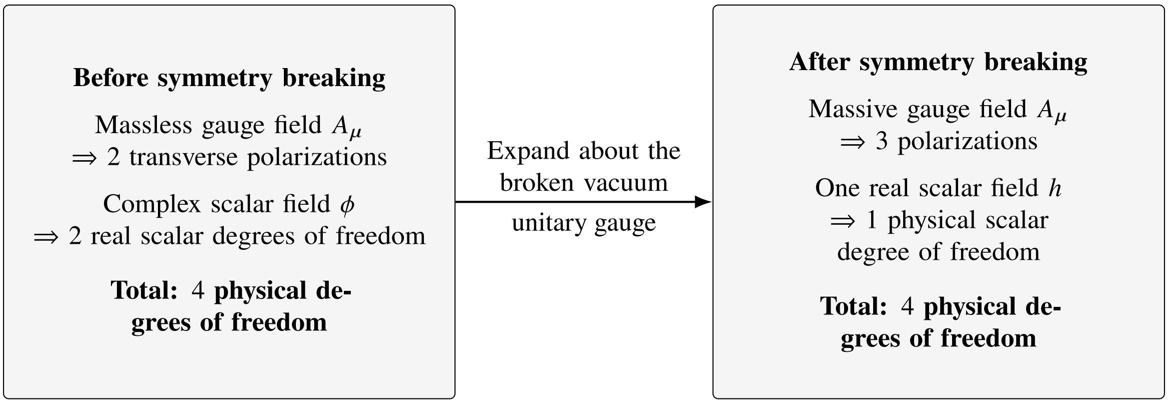

The would-be Goldstone mode does not remain as a physical massless scalar;
it becomes the longitudinal polarization of the massive vector boson.

Figure 5: Prototype form of the Higgs mechanism in a gauged $U(1)$ theory. The would-be Goldstone mode does not remain in the physical spectrum; instead, it is reinterpreted as the longitudinal polarization of the now massive vector boson, leaving one physical scalar excitation.

Remark 5.1: "The Goldstone boson is eaten" is shorthand

The phrase "the Goldstone boson is eaten" is pedagogically useful, but it should not be taken literally. No degree of freedom disappears. Rather, in the gauged theory the would-be Goldstone degree of freedom is reinterpreted as the longitudinal polarization of the now massive vector boson. The total number of physical degrees of freedom is unchanged.

5.5 Non-Abelian preview and the possibility of surviving unbroken gauge symmetry

In a non-Abelian gauge theory the same general mechanism operates, but the number of gauge bosons that become massive depends on the pattern of symmetry breaking. Gauge fields associated with broken generators acquire masses, while gauge fields associated with unbroken generators remain massless. This is precisely what happens in the electroweak Standard Model: the original symmetry $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$ is reduced to $\mathrm{U}(1)_{\mathrm{em}}$, so three gauge-field combinations become massive and one combination remains massless. The physical photon is the gauge field of the surviving unbroken electromagnetic symmetry. With this prototype mechanism in hand, we can now turn to its electroweak realization in the Standard Model itself.

6 The electroweak gauge structure before symmetry breaking

6.1 The gauge group $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$

The electroweak sector of the Standard Model is based on the gauge group

$$\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y. \tag{6.1}$$

The $\mathrm{SU}(2)_L$ factor acts on left-handed weak doublets, while the hypercharge group $\mathrm{U}(1)_Y$ distinguishes fields through the quantum number Y. Electromagnetism is not inserted as a separate gauge factor at this stage. Rather, it emerges as the unbroken subgroup after electroweak symmetry breaking.

6.2 Electroweak gauge fields and couplings

The gauge fields of $\mathrm{SU}(2)_L$ are three vector fields W_μ^a, $a = 1, 2, 3$, with coupling g. The gauge field of $\mathrm{U}(1)_Y$ is B_μ, with coupling g'. The electroweak covariant derivative acting on a field of weak-isospin generators $T^a = \tau^a/2$ and hypercharge Y is

$$D_\mu = \partial_\mu - \mathrm{i}gT^aW_\mu^a - \mathrm{i}g'YB_\mu. \tag{6.2}$$

The fact that there are initially four massless electroweak gauge fields is important. The observed low-energy spectrum contains one massless vector boson, the photon, and three heavy weak bosons. The Higgs mechanism must explain how this reorganisation occurs.

6.3 Left-handed doublets, right-handed singlets, and chirality

For one generation, the left-handed fermions are organised into weak doublets,

$$L_L = \begin{pmatrix} \nu_e \\ e \end{pmatrix}_L, \qquad Q_L = \begin{pmatrix} u \\ d \end{pmatrix}_L, \tag{6.3}$$

whereas the right-handed charged lepton and quarks are weak singlets,

$$e_R, \qquad u_R, \qquad d_R. \tag{6.4}$$

The three-generation theory repeats this pattern for the muon, tau, charm, strange, top, and bottom fields. The important point for the present module is not the multiplicity of generations but the chiral separation: left-handed and right-handed fields transform differently under the electroweak gauge group.

6.4 Why the electroweak field content already anticipates the mass problem

Once the field content is written this way, the mass problem is already present before any scalar field is introduced. Gauge invariance forbids direct masses for the gauge bosons. Chirality forbids naive Dirac masses for fermions. In this sense the Higgs sector is not an optional decoration attached later. It is the minimal scalar structure needed to convert the electroweak gauge theory into the observed low-energy massive theory without abandoning its symmetry principles.

7 The Higgs doublet and the scalar potential

7.1 Why the minimal scalar sector is a complex doublet

The minimal Standard Model introduces one complex scalar doublet. This is the smallest scalar representation that can break $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$ down to $\mathrm{U}(1)_{\mathrm{em}}$ while leaving one neutral scalar excitation in the spectrum. A smaller choice, such as a real singlet, would not couple correctly to the electroweak gauge fields and would not solve the weak-boson mass problem. A larger scalar sector is possible, but it would go beyond the minimal Standard Model.

7.2 Higgs quantum numbers and field content

The Higgs field is written as

$$\Phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix}, \qquad Y_\Phi = +\tfrac{1}{2}. \tag{7.1}$$

Because the doublet is complex, it contains four real scalar degrees of freedom. Before symmetry breaking those four degrees of freedom are simply the components of the electroweak scalar field. After symmetry breaking, three are reinterpreted as the longitudinal modes of the massive weak bosons and one remains as the physical Higgs boson.

7.3 Higgs kinetic term and electroweak covariant derivative

The electroweak Higgs sector contains the gauge-covariant kinetic term

$$\mathcal{L}_{\text{Higgs}} \supset (D_\mu \Phi)^\dagger (D^\mu \Phi), \tag{7.2}$$

with

$$D_\mu = \partial_\mu - \mathrm{i}g\frac{\tau^a}{2}W_\mu^a - \mathrm{i}g' Y B_\mu. \tag{7.3}$$

This term is structurally central. It tells us how the Higgs field transforms under the electroweak gauge symmetry, and after symmetry breaking it produces the mass terms of the gauge bosons and the Higgs–gauge-boson interaction terms.

7.4 The Higgs potential

The renormalisable scalar potential of the minimal Standard Model is

$$V(\Phi) = \mu^2 \Phi^\dagger \Phi + \lambda (\Phi^\dagger \Phi)^2, \tag{7.4}$$

with $\lambda > 0$ to ensure stability at large field values. The sign of μ^2 determines the qualitative shape of the potential. If $\mu^2 > 0$, the minimum is at $\Phi = 0$ and the symmetry is manifest in the vacuum. If $\mu^2 < 0$, the origin is unstable and the vacuum occurs at nonzero $\Phi^\dagger\Phi$.

7.5 Conditions on μ^2 and λ

The phenomenologically relevant case is

$$\mu^2 < 0, \qquad \lambda > 0. \tag{7.5}$$

Then the potential has the familiar "Mexican hat" form in field space. The quartic coupling keeps the potential bounded from below, while the negative quadratic term pushes the vacuum away from the origin. This is the electroweak realization of spontaneous symmetry breaking.

8 Vacuum expectation value and electroweak symmetry breaking

8.1 Minimisation of the Higgs potential

To find the vacuum one minimises $V(\Phi)$. Since the potential depends only on the gauge-invariant combination $\Phi^\dagger\Phi$, define

$$\rho^2 \equiv 2\Phi^\dagger\Phi. \tag{8.1}$$

Then

$$V = \tfrac{1}{2}\mu^2\rho^2 + \frac{\lambda}{4}\rho^4. \tag{8.2}$$

The minimum satisfies

$$\frac{\mathrm{d}V}{\mathrm{d}\rho} = \mu^2\rho + \lambda\rho^3 = 0, \tag{8.3}$$

so either $\rho = 0$ or

$$\rho^2 = -\frac{\mu^2}{\lambda} \equiv v^2. \tag{8.4}$$

Thus the nontrivial vacuum expectation value is

$$v = \sqrt{-\frac{\mu^2}{\lambda}}. \tag{8.5}$$

Definition 8.1: Vacuum expectation value

The vacuum expectation value (vev) of a field is its value in the ground state of the theory. In practical field-theory language, it is the field value around which one expands to describe physical excitations. A nonzero vev can signal spontaneous symmetry breaking when the vacuum is not invariant under the full symmetry of the Lagrangian.

8.2 Choice of vacuum direction

Because all points related by electroweak gauge transformations are physically equivalent, once the nonzero minimum has been established one may choose the vacuum in a convenient direction in the Higgs doublet by a gauge transformation. The standard choice is

$$\langle\Phi\rangle = \frac{1}{\sqrt{2}}\begin{pmatrix}0\\ v\end{pmatrix}. \tag{8.6}$$

This does not mean that the lower component is intrinsically special before the vacuum is chosen. Rather, it is a convenient gauge choice used to represent the broken vacuum.

8.3 The breaking pattern $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y \to \mathrm{U}(1)_{\mathrm{em}}$

The vacuum (8.6) is not invariant under the full electroweak group. However, it is invariant under a particular linear combination of generators, namely the electric charge operator

$$Q = T_3 + Y. \tag{8.7}$$

Therefore the symmetry-breaking pattern is

$$\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y \longrightarrow \mathrm{U}(1)_{\mathrm{em}}. \tag{8.8}$$

This is the crucial structural statement of electroweak symmetry breaking. Three generators are broken and one survives. The surviving unbroken symmetry is electromagnetism.

Example 8.1: Checking that the vacuum is neutral under $Q = T_3 + Y$

For the Higgs doublet with $Y = \frac{1}{2}$, the weak-isospin generator is

$$T_3 = \frac{1}{2}\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

Acting on the chosen vacuum

$$\langle \Phi \rangle = \frac{1}{\sqrt{2}}\begin{pmatrix} 0 \\ v \end{pmatrix},$$

one finds

$$Q\langle \Phi \rangle = (T_3 + Y)\langle \Phi \rangle = \frac{1}{\sqrt{2}}\begin{pmatrix} 0 \\ \left(-\frac{1}{2} + \frac{1}{2}\right) v \end{pmatrix} = 0.$$

Thus the vacuum is neutral under Q, which is why the electromagnetic subgroup remains unbroken.

8.4 The electroweak scale v

The parameter v is the electroweak vacuum expectation value. In the Standard Model it is related directly to the weak-boson masses and hence to the scale at which electroweak symmetry breaking becomes visible. Numerically, one finds $v \approx 246\,\text{GeV}$ from measured electroweak parameters. Pedagogically, the important point is not the number itself but its meaning: it is the size of the Higgs-field vacuum value around which the physical electroweak spectrum is organised.

8.5 Degree-of-freedom counting before and after symmetry breaking

Before symmetry breaking, the Higgs doublet contains four real scalar degrees of freedom. After symmetry breaking, three broken electroweak generators imply three would-be Goldstone modes. These are absorbed into the longitudinal polarization states of the massive W^+, W^-, and Z bosons. The remaining scalar degree of freedom appears as the physical Higgs boson h.

Thus the degree-of-freedom count is perfectly balanced:

- before breaking: four real scalar degrees of freedom and four massless electroweak gauge fields;
- after breaking: one physical scalar, one massless photon, and three massive weak bosons.

Counting physical degrees of freedom explicitly, one has 4 real scalar modes plus $4 \times 2 = 8$ transverse gauge-boson polarizations before symmetry breaking, and 1 scalar mode plus 2 photon polarizations plus $3 \times 3 = 9$ massive weak-boson polarizations after symmetry breaking, giving 12 physical degrees of freedom in both descriptions. This count is not merely a bookkeeping exercise. It is one of the clearest ways to see that the Higgs mechanism reorganises the theory consistently.

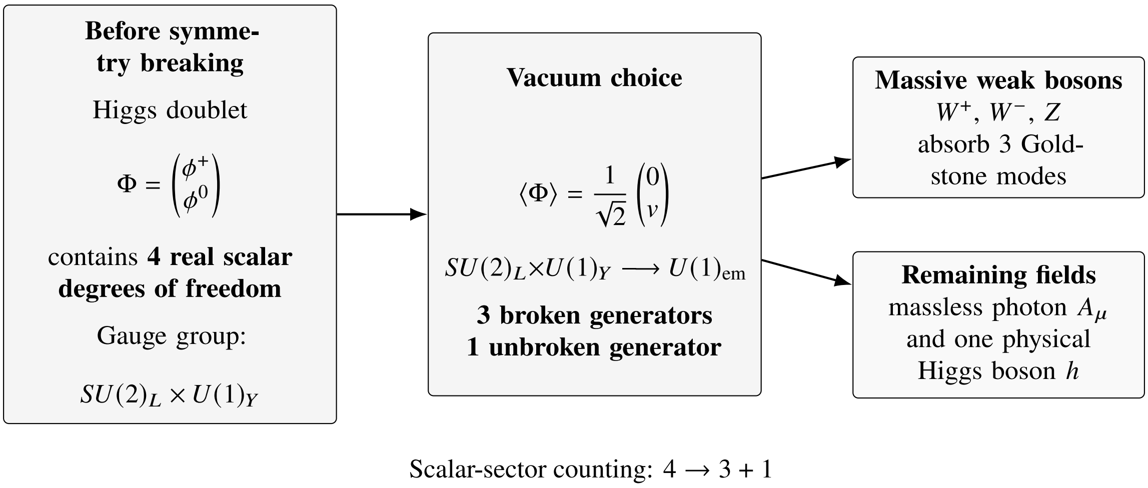

Figure 6: Electroweak symmetry breaking in the minimal Higgs-doublet theory. A complex scalar doublet contains four real degrees of freedom; after $SU(2)_L \times U(1)_Y \rightarrow U(1)_{\text{em}}$, three become the longitudinal modes of $W^\pm$ and Z, while one remains as the physical Higgs boson.

9 Gauge fixing, unitary gauge, and physical field content

9.1 Goldstone directions versus physical scalar direction

It is helpful to separate the Higgs fluctuations into one radial direction and three directions associated with the broken generators. The radial direction changes the length of the Higgs field in field space and corresponds to the physical Higgs boson. The other three directions move the field along the vacuum manifold and correspond to would-be Goldstone modes.

9.2 Unitary-gauge form of the Higgs doublet

In unitary gauge the three Goldstone directions are gauged away, and the Higgs doublet takes the simple form

$$\Phi(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h(x) \end{pmatrix}. \tag{9.1}$$

This is an especially convenient gauge for identifying the physical spectrum. The scalar field $h(x)$ is the remaining physical excitation around the electroweak vacuum.

Definition 9.1: Unitary gauge

Unitary gauge is a gauge choice in which the would-be Goldstone fields are removed from the scalar multiplet, so that the remaining physical field content is displayed explicitly. In the electroweak theory this makes the Higgs doublet take the form

$$\Phi(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h(x) \end{pmatrix},$$

with one visible physical scalar field $h(x)$.

9.3 Why three scalar modes are absorbed

The number three is dictated by the symmetry-breaking pattern. Since $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$ has four generators and $\mathrm{U}(1)_{\mathrm{em}}$ has one, the number of broken generators is three. The corresponding three would-be Goldstone modes are absorbed by the gauge fields associated with the broken generators. The resulting massive vector bosons require precisely three extra longitudinal modes: one for each of W^+, W^-, and Z.

9.4 One physical scalar remains

After the absorption of the three would-be Goldstone modes, one real scalar degree of freedom remains. This is the Higgs boson. It is therefore not an arbitrary additional particle. It is the unavoidable leftover scalar excitation once the electroweak gauge symmetry has been spontaneously broken by a single complex doublet.

10 Gauge-boson masses from the Higgs kinetic term

10.1 The electroweak covariant derivative in component form

To derive the gauge-boson masses, insert the unitary-gauge field (9.1) into the Higgs kinetic term. The covariant derivative acting on the vacuum part of the Higgs field is

$$D_\mu \langle \Phi \rangle = -\frac{\mathrm{i}}{\sqrt{2}} \left[g \frac{\tau^a}{2} W^a_\mu + g' Y B_\mu \right] \begin{pmatrix} 0 \\ v \end{pmatrix}. \tag{10.1}$$

Because the derivative acts on a constant vacuum expectation value, only the gauge fields survive. Squaring this term will therefore generate the gauge-boson mass matrix.

10.2 The charged gauge sector and the $W^\pm$ fields

The charged combinations are defined by

$$W^\pm_\mu = \frac{1}{\sqrt{2}} \left(W^1_\mu \mp \mathrm{i} W^2_\mu \right). \tag{10.2}$$

Inserting the Higgs vacuum expectation value into $(D_\mu\Phi)^\dagger(D^\mu\Phi)$ and isolating the W^1 and W^2 pieces first gives

$$\mathcal{L} \supset \frac{g^2v^2}{8}\left[(W^1_\mu)^2 + (W^2_\mu)^2\right]. \tag{10.3}$$

Using the definition of $W^\pm_\mu$, this may be rewritten as

$$\mathcal{L} \supset \frac{g^2v^2}{4} W^+_\mu W^{-\mu}. \tag{10.4}$$

10.3 Derivation of m_W

Comparing with the standard massive-vector normalization, one identifies

$$m_W = \frac{gv}{2}. \tag{10.5}$$

This equation is one of the central results of the module. It shows that the weak-boson mass is not inserted by hand; it follows from the electroweak gauge coupling and the Higgs vacuum expectation value.

Example 10.1: Reading off m_W from the Higgs kinetic term

Once the Higgs vacuum expectation value is inserted, the charged sector contains

$$\mathcal{L} \supset \frac{g^2v^2}{4} W^+_\mu W^{-\mu}.$$

Comparing this with the standard mass term for the charged vector field, one reads off

$$m_W^2 = \frac{g^2v^2}{4}, \qquad \text{so} \qquad m_W = \frac{gv}{2}.$$

The important point is structural: the weak-boson mass emerges from the gauge-invariant Higgs kinetic term rather than from an explicit mass term inserted by hand.

10.4 Longitudinal polarisation and the eaten Goldstone modes

A massless spin-1 particle has only two transverse physical polarizations. A massive spin-1 particle has three. The extra polarization of each massive weak boson is supplied by one of the absorbed Goldstone modes. The phrase "eaten Goldstone boson" is pedagogically vivid, but one should interpret it carefully: no degree of freedom disappears. Rather, a scalar degree of freedom is reinterpreted as the longitudinal polarization state of a massive vector boson.

11 The neutral electroweak sector and photon–Z mixing

11.1 The neutral-gauge-boson mass matrix

The neutral part of the Higgs kinetic term involves the fields W^3_μ and B_μ. Using $Y = +\frac{1}{2}$ for the Higgs doublet, one finds the quadratic neutral-field contribution

$$\mathcal{L}_{\text{neutral mass}} = \frac{v^2}{8} \begin{pmatrix} W^3_\mu & B_\mu \end{pmatrix} \begin{pmatrix} g^2 & -gg' \\ -gg' & g'^2 \end{pmatrix} \begin{pmatrix} W^{3\mu} \\ B^\mu \end{pmatrix}. \tag{11.1}$$

This matrix makes two important facts immediately visible. First, the neutral fields mix. Second, the determinant vanishes, so one eigenvalue is zero. That zero eigenvalue will correspond to the massless photon.

11.2 The Weinberg rotation

Diagonalising (11.1) leads to the physical neutral fields

$$A_\mu = \sin\theta_W\, W^3_\mu + \cos\theta_W\, B_\mu, \qquad Z_\mu = \cos\theta_W\, W^3_\mu - \sin\theta_W\, B_\mu, \tag{11.2}$$

where the weak mixing angle is defined by

$$\tan\theta_W = \frac{g'}{g}. \tag{11.3}$$

Definition 11.1: Weinberg rotation

The Weinberg rotation is the orthogonal change of basis that takes the neutral gauge eigenstates W^3_μ and B_μ into the physical neutral fields A_μ and Z_μ. It diagonalises the neutral electroweak mass matrix after symmetry breaking, and its rotation angle is the weak mixing angle θ_W.

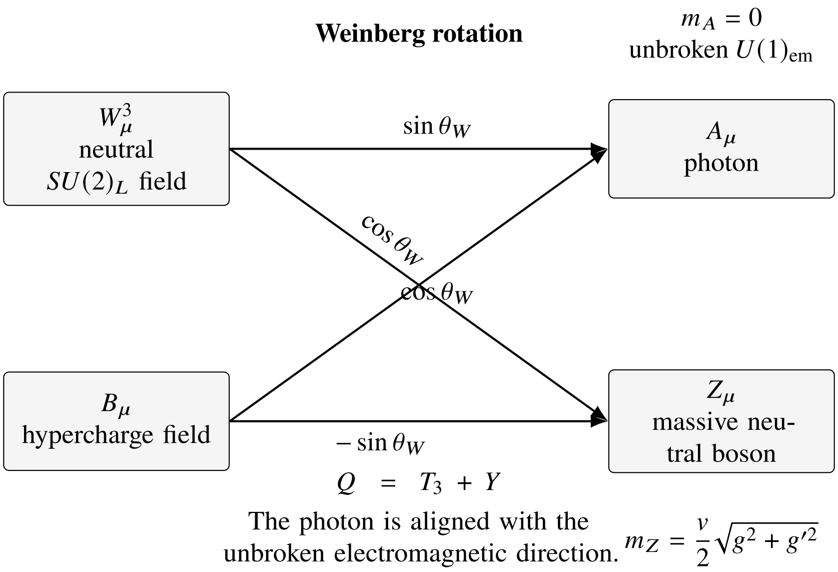

Figure 7: Neutral electroweak mixing after symmetry breaking. The original neutral gauge fields W^3_μ and B_μ are rotated into the physical photon and Z boson. The photon corresponds to the unbroken electromagnetic direction and therefore remains massless.

This transformation is called the Weinberg rotation. It rotates the original gauge eigenstates into the physical neutral bosons. Equivalently, the zero-eigenvalue direction of the matrix in (11.1) is proportional to (g', g), which is precisely the combination identified with the massless photon, while the orthogonal direction becomes the massive Z boson.

11.3 Derivation of m_Z and the massless photon

The nonzero eigenvalue of the neutral mass matrix gives the Z-boson mass,

$$m_Z = \frac{v}{2}\sqrt{g^2 + g'^2}, \tag{11.4}$$

while the orthogonal combination corresponding to A_μ remains massless,

$$m_A = 0. \tag{11.5}$$

Together with (10.5), this implies the well-known tree-level relation

$$\frac{m_W}{m_Z} = \cos\theta_W. \tag{11.6}$$

The photon is therefore not introduced separately. It emerges as the gauge boson of the unbroken subgroup.

11.4 The Weinberg angle

The angle θ_W is not a decorative parameter. It encodes the relative admixture of the original $\mathrm{SU}(2)_L$ and $\mathrm{U}(1)_Y$ gauge fields in the physical neutral bosons. It also determines the relation between the electroweak couplings and the electric charge,

$$e = g\sin\theta_W = g'\cos\theta_W. \tag{11.7}$$

Thus the electric coupling is not independent of the electroweak sector; it is the coupling of the surviving unbroken $\mathrm{U}(1)_{\mathrm{em}}$ subgroup.

11.5 The unbroken generator $Q = T_3 + Y$

Why does electromagnetism survive? The reason is encoded in the Higgs vacuum. Acting with the electric charge operator on the chosen vacuum gives zero:

$$Q\langle\Phi\rangle = (T_3 + Y)\langle\Phi\rangle = 0. \tag{11.8}$$

Therefore the vacuum is neutral under Q. The subgroup generated by Q is unbroken, while the other three combinations of generators are broken. This compact equation captures the full structure of electroweak symmetry breaking.

Remark 11.1: Why the photon remains massless

The photon is not massless because one simply chose to leave one eigenvalue equal to zero after diagonalising the neutral sector. It is massless because the Higgs vacuum is neutral under the generator $Q = T_3 + Y$. The surviving unbroken $U(1)_{\mathrm{em}}$ therefore has a massless gauge boson, namely the photon.

11.6 Physical interpretation of photon–Z mixing

The physical photon and Z boson are not the original fields of the Lagrangian before symmetry breaking. They are mixtures of the neutral gauge fields W^3_μ and B_μ. The photon corresponds to the unbroken direction in gauge space and so remains massless. The Z corresponds to the orthogonal broken direction and therefore becomes massive. Photon–Z mixing is thus not an optional phenomenological detail; it is the physical expression of how $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$ reorganises itself into a massive weak sector plus unbroken electromagnetism.

12 Why fermion masses require Yukawa interactions

12.1 Chiral structure revisited

The chiral assignments of the electroweak theory were not chosen arbitrarily. They encode the experimentally established fact that the weak interaction acts only on left-handed doublets. This chiral structure is central to the Standard Model, but it also forbids direct fermion mass terms. The theory therefore needs a new kind of interaction that couples left- and right-handed fields without violating gauge symmetry.

12.2 Why naive Dirac mass terms violate electroweak gauge symmetry

As emphasized earlier, a Dirac mass term couples left- and right-handed fermions. For the electron,

$$-m_e \bar{e} e = -m_e(\bar{e}_L e_R + \bar{e}_R e_L). \tag{12.1}$$

However, e_L is contained inside a weak doublet and e_R is a weak singlet. Since these pieces transform differently under $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$, the mass term is not gauge invariant. The same logic applies to quarks. This is not a defect of the formalism; it is the theory telling us that masses must arise through a gauge-invariant mediator.

12.3 Gauge-invariant scalar–fermion couplings as the cure

The cure is to couple the fermions to the Higgs doublet. Since the Higgs field is itself charged under the electroweak group, one can combine left-handed doublets, right-handed singlets, and the Higgs field into gauge-invariant products. These are the Yukawa interactions. When the Higgs field is replaced by its vacuum expectation value plus its fluctuation, these Yukawa terms become both fermion mass terms and Higgs–fermion interaction terms.

Definition 12.1: Yukawa interaction

A Yukawa interaction is a gauge-invariant coupling between fermion fields and a scalar field, linear in the scalar and bilinear in the fermions. In the Standard Model, Yukawa terms connect left-handed electroweak doublets to right-handed singlets through Φ or $\tilde{\Phi}$. After electroweak symmetry breaking, these same terms generate both fermion masses and Higgs–fermion couplings.

12.4 What this solves and what it does not solve

Yukawa interactions explain the *mechanism* by which fermion masses arise in the Standard Model. They do not explain why the Yukawa couplings take the numerical values they do. The mass hierarchy of the fermions, and the broader flavour pattern of the Standard Model, remain among the unresolved features of the theory. This important limitation will reappear later when we discuss what the Higgs sector does and does not explain.

13 Yukawa interactions and fermion mass generation

13.1 Lepton Yukawa terms

For one generation, the charged-lepton Yukawa term may be written as

$$\mathcal{L}_{\text{Yuk}}^{(\ell)} = -y_e\, \bar{L}_L \Phi e_R + \text{h.c.}. \tag{13.1}$$

This term is gauge invariant because the electroweak quantum numbers of L_L, Φ, and e_R combine to form a singlet.

13.2 Down-type quark Yukawa terms

The down-type quark Yukawa term has the analogous form

$$\mathcal{L}_{\text{Yuk}}^{(d)} = -y_d\, \bar{Q}_L \Phi d_R + \text{h.c.}. \tag{13.2}$$

Again, the Higgs doublet is the bridge that connects the weak doublet to the weak singlet in a gauge-invariant way.

13.3 Up-type quark Yukawa terms and the conjugate Higgs doublet

For up-type quarks one needs the conjugate Higgs doublet, conventionally denoted by $\tilde{\Phi}$,

$$\widetilde{\Phi} \equiv \mathrm{i}\tau^2 \Phi^*. \tag{13.3}$$

Its hypercharge is opposite to that of Φ, and the Yukawa term is

$$\mathcal{L}_{\text{Yuk}}^{(u)} = -y_u\, \bar{Q}_L \widetilde{\Phi} u_R + \text{h.c.}. \tag{13.4}$$

This is one of the places where the specific electroweak quantum numbers of the Higgs doublet matter decisively.

13.4 Fermion masses after electroweak symmetry breaking

Insert the unitary-gauge Higgs field

$$\Phi(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h(x) \end{pmatrix} \tag{13.5}$$

into the charged-lepton Yukawa term. One obtains

$$-y_e \bar{L}_L \Phi e_R + \text{h.c.} \longrightarrow -\frac{y_e}{\sqrt{2}}(v+h)\bar{e}e. \tag{13.6}$$

This separates naturally into a mass term and an interaction term,

$$-\frac{y_e v}{\sqrt{2}}\bar{e}e - \frac{y_e}{\sqrt{2}}h\bar{e}e. \tag{13.7}$$

Exactly the same logic applies to quarks.

Example 13.1: One Yukawa term gives both mass and coupling

For the charged lepton,

$$-\, y_e\, \bar{L}_L \Phi e_R + \text{h.c.} \quad \longrightarrow \quad -\frac{y_e}{\sqrt{2}}(v+h)\,\bar{e}e.$$

This immediately separates into

$$-\frac{y_e v}{\sqrt{2}}\,\bar{e}e \; - \; \frac{y_e}{\sqrt{2}}\,h\,\bar{e}e.$$

The first term is the fermion mass term, so $m_e = y_e v/\sqrt{2}$, while the second is the Higgs–fermion interaction, which may be rewritten as $-(m_e/v)\,h\,\bar{e}e$.

13.5 $m_f = y_f v/\sqrt{2}$

For any charged fermion in the minimal Standard Model one obtains the relation

$$m_f = \frac{y_f v}{\sqrt{2}}. \tag{13.8}$$

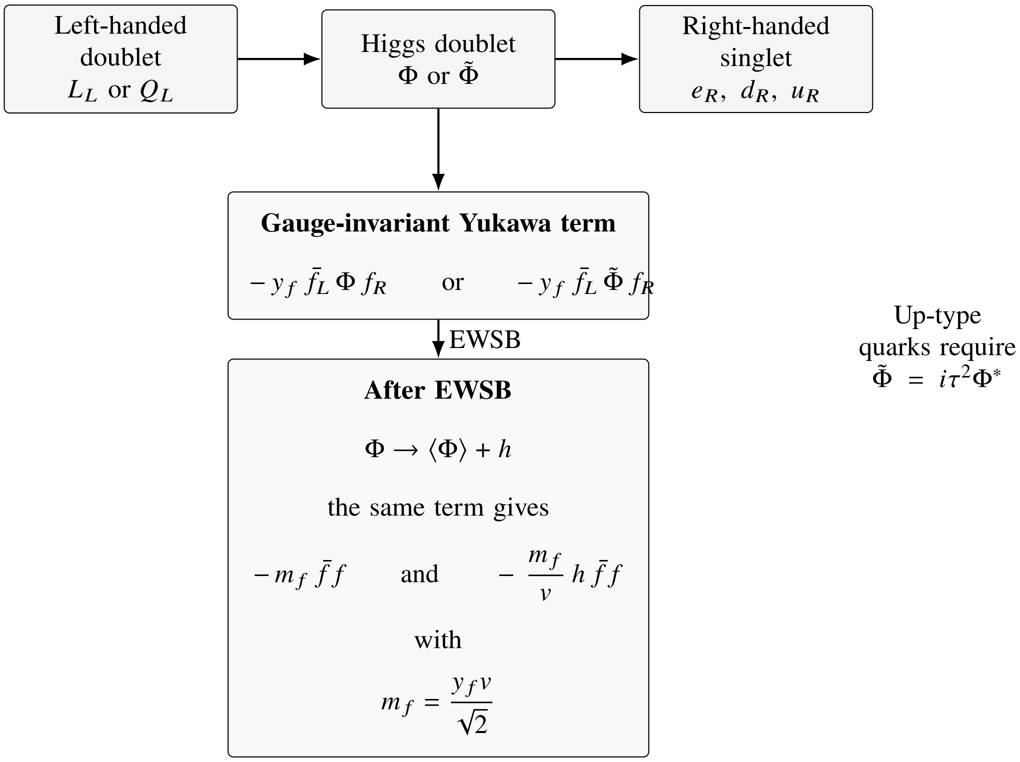

Figure 8: Gauge-invariant Yukawa couplings connect left-handed electroweak doublets to right-handed singlets through the Higgs doublet. After electroweak symmetry breaking, the same interaction produces both the fermion mass term and the Higgs–fermion coupling.

This is the central fermion-mass formula of the module. It shows that the fermion masses are proportional to the Higgs vacuum expectation value and to the corresponding Yukawa coupling. The common scale v is shared by all fermions; the differences in their masses are encoded in the Yukawa couplings themselves.

13.6 Higgs–fermion couplings and proportionality to mass

The interaction term obtained above can be rewritten using (13.8) as

$$\mathcal{L}_{hff} = -\frac{m_f}{v}\,h\bar{f}f. \tag{13.9}$$

This simple equation has major physical consequences. It explains why the Higgs boson couples most strongly to the heaviest accessible fermions. It also shows why Higgs measurements directly probe the same structure that generates fermion masses.

13.7 Short orientation remark on flavour structure

In the realistic three-generation Standard Model the Yukawa couplings are matrices in flavour space, and diagonalising them leads to the familiar flavour structure of the quark and lepton sectors. Those details, including the CKM matrix and the mass eigenbasis, belong to a later flavour discussion. For the purposes of Module 5, the essential point is already visible in the one-generation form: fermion masses arise from gauge-invariant Yukawa couplings to the Higgs field.

14 The physical Higgs boson in the Standard Model

14.1 The Higgs boson as the surviving scalar excitation

After electroweak symmetry breaking and gauge fixing, the Higgs doublet contains one physical scalar field,

$$\Phi(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h(x) \end{pmatrix}. \tag{14.1}$$

The field $h(x)$ is the Higgs boson. It should be viewed as the scalar fluctuation of the electroweak vacuum itself. This interpretation is much deeper than saying that the Higgs is simply another particle in the Standard Model chart. It is the remaining degree of freedom of the scalar field whose vacuum expectation value reorganised the electroweak theory.

Definition 14.1: Physical Higgs boson

The physical Higgs boson is the remaining neutral scalar excitation after the three would-be Goldstone modes of the Higgs doublet have been absorbed into the longitudinal modes of $W^\pm$ and Z. In the minimal Standard Model it is the fluctuation field $h(x)$ around the electroweak vacuum.

14.2 Higgs mass from the scalar potential

Insert the unitary-gauge form of the Higgs field into the potential (7.4). Since

$$\Phi^\dagger \Phi = \frac{(v+h)^2}{2}, \tag{14.2}$$

one obtains

$$V(h) = \frac{\mu^2}{2}(v+h)^2 + \frac{\lambda}{4}(v+h)^4. \tag{14.3}$$

Using the minimisation condition $\mu^2 = -\lambda v^2$, the linear term in h cancels, and the quadratic term becomes

$$V(h) \supset \lambda v^2 h^2 = \tfrac{1}{2}(2\lambda v^2)h^2. \tag{14.4}$$

Therefore the Higgs mass is

$$m_h^2 = 2\lambda v^2. \tag{14.5}$$

The parameter λ therefore controls the Higgs mass once the vacuum expectation value is fixed.

14.3 Higgs self-interactions at orientation level

The same expansion of the potential also produces cubic and quartic self-interaction terms,

$$V(h) = \text{const} + \tfrac{1}{2}m_h^2 h^2 + \frac{m_h^2}{2v}h^3 + \frac{m_h^2}{8v^2}h^4. \tag{14.6}$$

These terms are not central to the present module, but they are worth displaying at orientation level because they show that the Higgs field interacts with itself. They arise from the same scalar potential that generated the nonzero vacuum expectation value, so they are another direct reflection of the structure of the Higgs sector. Such self-interactions are among the more difficult quantities to probe experimentally.

14.4 Why the Higgs is special in the Standard Model

The Higgs boson is special in at least three ways. First, it is the only fundamental scalar discovered so far. Second, its vacuum expectation value is responsible for the observed mass pattern of weak bosons and charged fermions. Third, the scalar sector contains some of the least explained parameters of the Standard Model, such as the Yukawa couplings and the quartic scalar coupling. The Higgs sector is therefore simultaneously one of the most successful and one of the most mysterious parts of the Standard Model.

15 Higgs couplings and basic observable properties

15.1 Higgs couplings to W and Z

The gauge-boson mass terms arose from the Higgs kinetic term after replacing Φ by its vacuum expectation value. If, instead, one keeps the fluctuation field $h(x)$ in the unitary-gauge expression, the same term generates interactions between one Higgs boson and two gauge bosons. Schematically,

$$(D_\mu \Phi)^\dagger (D^\mu \Phi) \supset m_W^2 W_\mu^+ W^{-\mu} \left(1 + \frac{h}{v}\right)^2 + \frac{m_Z^2}{2} Z_\mu Z^\mu \left(1 + \frac{h}{v}\right)^2 . \tag{15.1}$$

This yields couplings such as

$$\mathcal{L} \supset \frac{2m_W^2}{v} h W_\mu^+ W^{-\mu} + \frac{m_Z^2}{v} h Z_\mu Z^\mu + \cdots . \tag{15.2}$$

The Higgs thus couples directly to the weak gauge bosons, and the coupling strengths are proportional to their masses.

15.2 Higgs couplings to fermions

The fermionic analogue was already derived in (13.9):

$$\mathcal{L}_{hff} = -\frac{m_f}{v} h \bar{f} f . \tag{15.3}$$

This explains the basic qualitative pattern of Higgs phenomenology. Heavy particles couple more strongly to the Higgs than light particles. The Higgs is therefore especially sensitive to the top quark, the weak bosons, the bottom quark, and the tau lepton.

15.3 Spin, charge, and scalar character

The Higgs boson is electrically neutral and carries spin zero. In the minimal Standard Model it is a CP-even scalar excitation of the broken electroweak vacuum. These properties distinguish it sharply from the spin-1 gauge bosons and the spin-1/2 matter fields. The scalar character of the Higgs is not a superficial label: it is the reason why a renormalisable potential for the field can exist and why the field can acquire a nonzero vacuum expectation value without breaking Lorentz invariance.

15.4 Why Higgs couplings track particle masses

Equations (15.2) and (13.9) show a common pattern. The Higgs couplings to Standard Model particles are tied directly to the masses those particles acquire after electroweak symmetry breaking. The Higgs is therefore not merely associated with mass generation historically; it remains dynamically coupled to the same fields whose masses arise from the broken electroweak vacuum. This is why precision measurements of Higgs couplings are such direct tests of the Standard Model mechanism itself.

15.5 A bounded first look at decay patterns

For a Higgs boson near 125 GeV, the dominant decay is into $b\bar{b}$, while important subleading channels include WW^*, ZZ^*, $\tau^+\tau^-$, gg, and $\gamma\gamma$. The channels involving gauge bosons are especially revealing pedagogically because they make the Higgs–vector-boson couplings directly visible. The $\gamma\gamma$ mode is loop-induced and has a small branching ratio, but it is experimentally clean and therefore played an outsized role in discovery. The $ZZ^* \to 4\ell$ mode is even cleaner, although rarer.

Decay mode	Branching ratio
$H \to b\bar{b}$	57.8%
$H \to WW^*$	21.6%
$H \to \tau^+\tau^-$	6.4%
$H \to gg$	8.6%
$H \to c\bar{c}$	2.9%
$H \to ZZ^*$	2.7%
$H \to \gamma\gamma$	0.2%

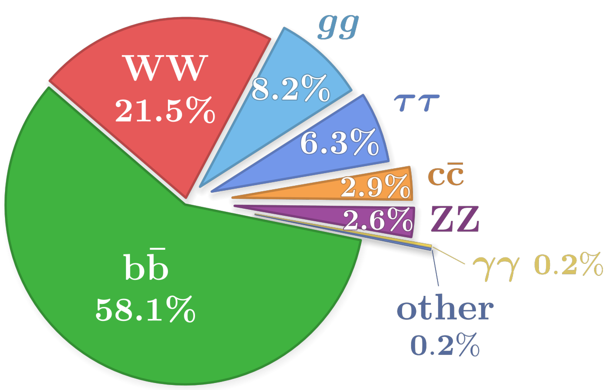

Figure 9: Representative Standard Model Higgs decay pattern near $m_H \simeq 125$ GeV. The dominant mode is $H \to b\bar{b}$, while WW^*, ZZ^*, $\tau^+\tau^-$, gg, and $\gamma\gamma$ remain especially important for interpretation and measurement.

15.6 What belongs here and what is deferred to Module 6

At this stage the aim is only to show how the Higgs couplings determine the broad pattern of decays and signatures. A full treatment of decay widths, branching ratios, production cross sections, and amplitude-level calculations belongs naturally to a phenomenology module. The present note therefore stops at the point where the qualitative link between electroweak structure and observable Higgs signatures has become clear.

Guided checks

- Can you explain why the same Higgs-sector terms that generate W and Z masses also imply hWW and hZZ couplings?
- Can you derive qualitatively why Higgs couplings to fermions are larger for heavier fermions?
- Can you state the difference between the physical Higgs boson and the three would-be Goldstone modes?
- Why can a decay channel with a small branching ratio still be experimentally important at a hadron collider?

16 Why the Higgs is experimentally accessible at the LHC

The preceding sections established the Higgs boson's structural role inside the Standard Model. We now add a bounded qualitative bridge showing why that scalar is experimentally accessible at the LHC.

16.1 Proton collisions, partons, and access to the electroweak scale

The LHC is a proton–proton collider, not an elementary-particle collider in the same sense as an e^+e^- machine. The hard interactions occur between partons — quarks and gluons — inside the incoming protons. Because the LHC operates at multi-TeV centre-of-mass energy and enormous luminosity, it has abundant access to the electroweak scale. This makes it possible to produce Higgs bosons even though the Higgs is neutral and not itself a constituent of the proton.

16.2 Why a hadron collider can produce a neutral scalar abundantly

A hadron collider is especially powerful when the relevant processes can draw on large parton luminosities. The proton contains a substantial gluon density over the range of momentum fractions relevant to Higgs production. Although the Higgs boson is electrically neutral and colourless, it can still be produced efficiently through partonic subprocesses that involve gluons, heavy quarks, and weak bosons. The collider question is therefore not whether the Higgs can couple directly to the proton as a whole, but whether the proton's partonic content can generate it through short-distance Standard Model interactions.

16.3 Role of gluons, heavy quarks, and electroweak bosons in Higgs production

The gluons matter because the dominant Higgs production mechanism at the LHC is gluon fusion through a heavy-quark loop, dominated by the top quark. The heavy top quark matters because the Higgs couples proportionally to mass. The electroweak bosons matter because they allow vector-boson fusion and associated production with W or Z bosons. Thus the same couplings that emerged from the Higgs mechanism also control the principal ways in which the Higgs is produced at hadron colliders.

16.4 Why production and decay must be discussed together

A Higgs boson is not observed directly as a stable particle. It is inferred from its production and from its decay products. For that reason, one cannot discuss Higgs production at the LHC in isolation from Higgs decay channels. Some production modes are common but experimentally messy, while some decay channels are rare but clean. For the purposes of this module, the point is only qualitative: collider visibility depends on the interplay of production mode, decay channel, and background environment.

17 Main Standard Model Higgs production mechanisms at the LHC

17.1 Gluon fusion (ggF)

The dominant Standard Model Higgs production mode at the LHC is gluon fusion,

$$gg \to H, \tag{17.1}$$

mediated at leading order by a loop of heavy quarks, predominantly the top quark. This process is pedagogically striking because the Higgs is colourless, yet gluons still dominate its production rate through quantum loops. It is therefore a vivid example of how QCD and electroweak physics are intertwined at hadron colliders.

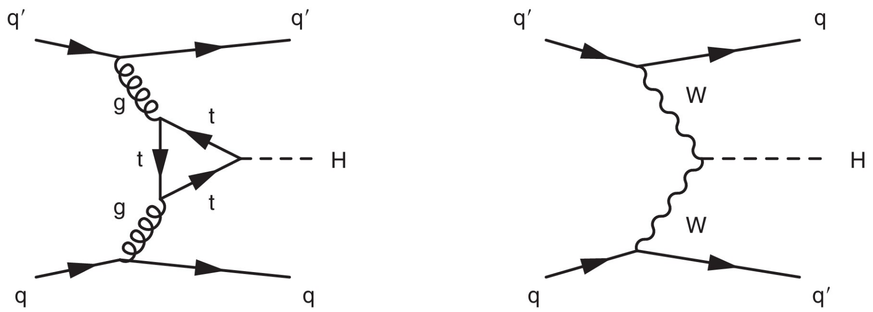

Figure 10: Two of the most important Standard Model Higgs production mechanisms at the LHC. Gluon fusion dominates the total production rate because of the large gluon luminosity in the proton and the strong coupling of the Higgs to the heavy top quark, while vector-boson fusion provides a cleaner topology and a more direct handle on Higgs couplings to weak bosons.

17.2 Why gluon fusion dominates

Two ideas explain the dominance of gluon fusion. First, the proton at LHC energies contains a very large gluon parton luminosity. Second, the Higgs couples most strongly to the heaviest Standard Model particles, so the top-quark loop gives the main contribution to the effective ggH interaction. The large gluon density compensates for the loop suppression and makes gluon fusion the largest production channel.

17.3 Vector-boson fusion (VBF)

In vector-boson fusion, two quarks from the incoming protons radiate weak bosons that then fuse into a Higgs boson. The final state typically contains the Higgs decay products together with two energetic forward jets. Although VBF is less abundant than ggF, its topology is often cleaner experimentally. This channel is especially useful pedagogically because it probes the Higgs couplings to weak bosons in a relatively direct way.

17.4 Associated production with weak bosons (VH)

In associated production, often called Higgs-strahlung, the Higgs is produced together with a W or Z boson:

$$q\bar{q}' \to WH, \qquad q\bar{q} \to ZH. \tag{17.2}$$

These channels are smaller than ggF and VBF, but they offer useful tagged environments. Leptonic decays of the associated weak boson can help isolate the event class, which is especially valuable for decay modes such as $H \to b\bar{b}$.

17.5 Associated production with top quarks (ttH)

In $t\bar{t}H$ production, the Higgs is created together with a top–antitop pair. This channel is rarer than ggF, VBF, or VH, but conceptually it is extremely important because it connects directly to the top Yukawa coupling. Since the top quark is the heaviest Standard Model fermion, its interaction with the Higgs is especially central to the mass-generation story.

17.6 Single-top-associated production (tH) as an advanced remark

Single-top-associated Higgs production is much rarer than the channels above, but it is theoretically interesting because different diagrams can interfere in a way that makes the process sensitive to the relative sign and size of the Higgs couplings involved. For the present module it is enough to know that tH exists as a more delicate subsidiary channel rather than as one of the main discovery modes.

17.7 Qualitative production hierarchy and complementary signatures

The main Standard Model Higgs production channels at the LHC may be summarized qualitatively as in Table 1. The essential lesson here is purely orienting: ggF is the largest mode, but different channels probe different couplings and offer different experimental environments.

Table 1: Qualitative pedagogical summary of the principal Standard Model Higgs production modes at the LHC.

Mode	Physical idea	Pedagogical role
ggF	Gluon fusion through a heavy-quark loop, dominated by the top quark	Largest rate; shows how gluons and the top Yukawa control Higgs production
VBF	Fusion of electroweak vector bosons radiated from quarks	Cleaner topology with forward jets; probes Higgs–vector-boson couplings
VH	Associated production with W or Z	Useful tagged environment; important for channels such as $H \to b\bar{b}$
ttH	Associated production with a top pair	Direct probe of the top Yukawa coupling
tH	Associated production with a single top quark	Advanced, rarer channel with useful interference sensitivity

18 Higgs decays, search channels, and discovery logic

18.1 Why $H \to \gamma\gamma$ and $H \to ZZ^* \to 4\ell$ were discovery channels

The two channels that played the most visible role in the original Higgs discovery were $H \to \gamma\gamma$ and $H \to ZZ^* \to 4\ell$. Neither channel is dominant in branching ratio, but both are exceptionally clean experimentally. Two energetic photons can be reconstructed with excellent invariant-mass resolution. Likewise, the four-lepton final state from ZZ^* decay is extraordinarily distinctive and allows precise event-by-event mass reconstruction.

The pedagogical importance of these channels is that they illustrate a general collider lesson: the most useful discovery channels are not always the largest decay modes. Small but clean signals can be more powerful than large but background-dominated ones.

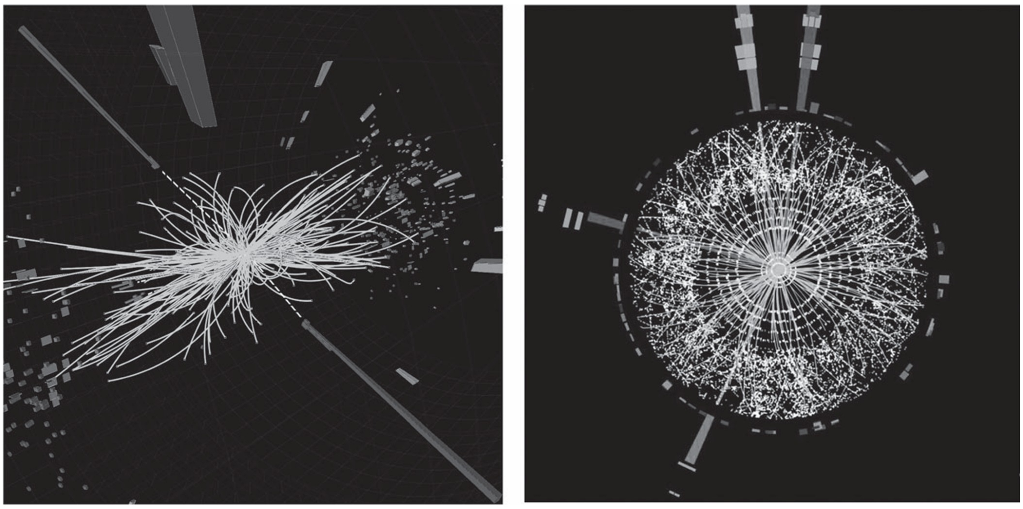

Figure 11: Representative Higgs-candidate events in the two cleanest discovery channels. The detector signatures illustrate why diphoton and four-lepton final states were exceptionally powerful in the early Higgs searches despite their relatively small branching fractions.

18.2 $H \to WW^*$, $H \to \tau\tau$, and $H \to b\bar{b}$ in the measurement programme

After discovery, additional decay channels became important in the broader Higgs measurement programme. The WW^* channel is important because of the large Higgs coupling to weak bosons. The $\tau^+\tau^-$ mode provides direct access to a fermionic Yukawa coupling in the lepton sector. The $b\bar{b}$ mode is the largest decay channel near 125 GeV but is difficult at hadron colliders because of the enormous QCD background. For this reason it is often discussed in more selective production environments such as VH.

18.3 Clean channels versus large backgrounds

Hadron colliders are rich in QCD activity. Multi-jet events, heavy-flavour production, and radiation from the incoming partons produce large backgrounds to many Higgs signatures. This is why channels with photons or multiple charged leptons were so important in discovery. They are not large in branching ratio, but they stand out above the background more clearly than generic hadronic final states. For Module 5, the key lesson is simply that the most informative channels are selected through a balance among branching fraction, cleanliness, and production environment.

18.4 Discovery in 2012 and later confirmation of the spin-0 Higgs boson

A new boson compatible with the Standard Model Higgs was announced by the ATLAS and CMS collaborations in 2012. The discovery established the existence of a new particle near 125 GeV. Subsequent studies supported an interpretation consistent with a spin-zero boson of the type expected for the Higgs sector of the Standard Model. In this sense the discovery did not merely add one more particle to the spectrum; it completed the minimal Standard Model particle content in a highly nontrivial way.

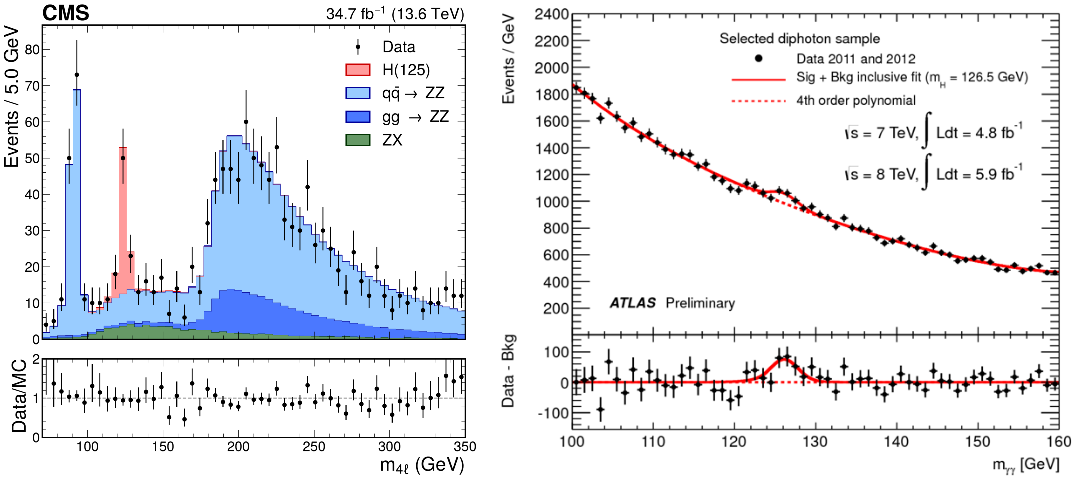

Figure 12: Discovery-era Higgs mass peaks in the two cleanest search channels. The diphoton and four-lepton invariant-mass distributions exhibit localized excesses near 125 GeV above the expected backgrounds, providing the most visually striking early evidence for the Higgs boson.

18.5 What the LHC has established about Higgs couplings

At the level needed here, LHC measurements support the qualitative picture that the observed Higgs boson is produced in the principal Standard Model channels and decays into the expected classes of final states in a way broadly compatible with Standard Model expectations. In particular, the data support the view that the Higgs couples to weak bosons and to fermions in a pattern tied to electroweak mass generation. For the purposes of this module, the main message is already clear: the observed Higgs behaves as the scalar excitation associated with electroweak symmetry breaking.

18.6 Pedagogical boundary: this is not a full experimental Higgs review

This section is deliberately bounded. It is meant to show why the Higgs boson is experimentally visible and why certain channels became discovery and measurement channels. It is not intended to replace a full collider-analysis treatment, nor to provide the machinery of detailed cross sections, likelihood fits, or detector-level reconstruction strategies. Those belong to a later phenomenology or laboratory discussion.

19 From Higgs Production Mechanisms to MadGraph Event Generation

The previous sections developed the Higgs sector from first principles: the mass problem of the electroweak theory, spontaneous symmetry breaking, the Higgs mechanism, electroweak symmetry breaking, Yukawa mass generation, and the physical Higgs boson. The module then connected this structure to the qualitative Higgs-production and discovery picture at the LHC. This section adds one bounded practical bridge. Its purpose is not to turn Module 5 into a full collider-analysis course, but to show how the Higgs couplings derived earlier become concrete partonic processes whose leading-order cross sections and event samples can be generated with MadGraph5_aMC@NLO [1].

19.1 Why a practical MadGraph bridge belongs here

The logic of the module has led to a simple but powerful result: the physical Higgs boson is not an arbitrary scalar added to the Standard Model spectrum. It is the surviving scalar excitation around the electroweak vacuum. The same electroweak symmetry-breaking mechanism that produces the masses of the weak bosons and charged fermions also fixes the leading pattern of Higgs interactions. In particular, the Higgs couplings to weak gauge bosons include

$$\mathcal{L} \supset \frac{2m_W^2}{v} hW^+_\mu W^{-\mu} + \frac{m_Z^2}{v} hZ_\mu Z^\mu, \tag{19.1}$$

and the Higgs coupling to a charged fermion is

$$\mathcal{L}_{hff} = -\frac{m_f}{v} h\bar{f}f. \tag{19.2}$$

These formulae are not only mass-generation statements. They are also collider statements. They tell us which particles can radiate, fuse into, or be produced together with a Higgs boson.

For example, the coupling to weak bosons is responsible for Higgs-strahlung and vector-boson-fusion processes, while the top Yukawa coupling controls associated production with top quarks and dominates the loop-induced gluon-fusion process. MadGraph provides a practical way to turn this coupling information into partonic amplitudes, cross sections, and event samples. The conceptual chain is

$$\mathcal{L}_{hVV},\ \mathcal{L}_{hff} \quad \Longrightarrow \quad \mathcal{M}(ab \to H + X) \quad \Longrightarrow \quad \sigma_{\text{LO}} \quad \Longrightarrow \quad \text{parton-level events.} \tag{19.3}$$

Remark 19.1: Notation: field versus external particle

In the formal parts of this module, the symbol $h(x)$ or h denotes the physical Higgs field appearing in the Lagrangian. In collider process notation it is conventional to write H for an external Higgs boson in the final state, as in $pp \to Hjj$ or $e^+e^- \to ZH$. In this section, h is used for Lagrangian-level formulae and H is used for collider processes.

The calculations discussed below should be read as leading-order orientation exercises. They are useful for learning which couplings control which production mechanisms, how collider initial states differ, and how event-generator workflows are organised. They should not be confused with state-of-the-art Higgs predictions. Precision Higgs cross sections require higher-order QCD and electroweak corrections, careful treatment of theoretical uncertainties, and comparison with dedicated references such as the LHC Higgs Cross Section Working Group reports [2, 3, 4].

19.2 Collider environments and coupling messages

Different colliders probe the same Higgs sector in complementary ways. A proton–proton collider has large parton luminosities and rich QCD radiation. An electron–positron collider has a clean and well-defined initial state. A lepton–proton collider combines a clean lepton beam with a partonic proton beam, making electroweak Higgs production in charged-current and neutral-current topologies especially transparent.

Collider environment	Initial state	Representative Higgs mechanisms	Main coupling message	Practical MadGraph issue
LHC	pp	ggF, VBF, VH, $t\bar{t}H$, tH	y_t, HWW, HZZ, Hff	proton PDFs, QCD radiation, scales, and loop-induced ggF
FCC-ee	e^+e^-	ZH, WW-fusion, ZZ-fusion	HZZ, HWW	clean fixed beams; no proton PDFs in the simplest exercises
LHeC-like	$e^{\pm}p$	CC $ep \to \nu Hj$, NC $ep \to eHj$	HWW, HZZ	one lepton beam plus one proton beam; PDF only on the proton side

Table 2: Complementary Higgs-production environments. The same Higgs couplings derived from electroweak symmetry breaking appear in different practical forms depending on the collider initial state.

At the LHC, the large gluon density inside the proton makes gluon fusion the dominant inclusive Higgs-production mechanism, even though the Higgs is colourless and the process is loop-induced. VBF and associated production are less abundant but often provide cleaner tags and more direct sensitivity to Higgs couplings to weak bosons. At FCC-ee, the process $e^+e^- \to ZH$ gives a particularly clean handle on the HZZ interaction, while WW-fusion probes HWW. In an LHeC-like environment, charged-current Higgs production is naturally associated with W-boson exchange and neutral-current production with Z-boson exchange, making the separation between HWW and HZZ topologies especially instructive [6].

19.3 What MadGraph calculates at this level

MadGraph reads a model file containing particles, parameters, and interactions, then generates the Feynman diagrams and matrix elements for a requested process. For the leading-order exercises considered here, the basic object is the tree-level or effective leading-order matrix element for a specified partonic scattering. For a fixed initial state ab, one may schematically write

$$\sigma_{\rm LO}(ab \to X) = \frac{1}{2\hat{s}} \int d\Phi_X \, \overline{|\mathcal{M}_{\rm LO}(ab \to X)|^2}, \tag{19.4}$$

where $\hat{s}$ is the partonic centre-of-mass energy squared, $d\Phi_X$ is the final-state phase-space measure, and the bar denotes the usual average over initial-state quantum numbers and sum over final-state quantum numbers.

For a proton–proton collider, the partonic calculation must be folded with parton distribution functions. The hadron-level cross section has the schematic factorised form

$$\sigma(pp \to X) = \sum_{i,j} \int_0^1 dx_1\, dx_2\, f_i(x_1, \mu_F) f_j(x_2, \mu_F) \hat{\sigma}_{ij}(x_1 x_2 s, \mu_F, \mu_R). \tag{19.5}$$

Here x_1 and x_2 are the momentum fractions carried by the incoming partons, f_i and f_j are PDFs evaluated at the factorisation scale μ_F, and μ_R is the renormalisation scale. This is why proton-beam simulations require a PDF interface such as LHAPDF [7].

For simple e^+e^- exercises, the initial state is fixed by the beam energy and no proton PDFs appear. For

$e^{\pm}p$ exercises, one beam is lepton-like and the other is proton-like: the lepton beam has no proton PDF, while the proton beam must be described by PDFs. This distinction is one of the most important practical lessons for students moving from formal Higgs couplings to collider simulations.

Take-home message

At the level of this module, MadGraph should be viewed as a bridge from Lagrangian interactions to leading-order parton-level collider processes. The output depends on the chosen model, beam energies, PDFs, scales, cuts, masses, widths, and perturbative order. A number produced by a simple MadGraph run is therefore not a universal "the Higgs cross section"; it is a prediction within a precisely specified setup.

19.4 Higgs production at the LHC in MadGraph

At a proton–proton collider, the incoming protons are not elementary partons in the hard scattering. The actual short-distance process is initiated by quarks, antiquarks, or gluons drawn from the PDFs. The main Standard Model Higgs-production mechanisms can be organised as follows.

Mode	Physics idea	Representative process	Coupling message	MadGraph comment
ggF	gluon fusion through a heavy-quark loop, dominated by the top quark	$gg \to H$, or $pp \to H$ in an effective model	top loop / effective ggH	loop-induced in the full SM; HEFT is useful but approximate
VBF	two quarks radiate weak bosons that fuse into a Higgs boson	$pp \to Hjj$	HWW, HZZ	distinguish electroweak VBF from generic QCD-induced Hjj
VH	associated production with a weak boson	$pp \to WH$, $pp \to ZH$	HVV	useful tagged environment, especially for selected Higgs decays
$t\bar{t}H$	Higgs production with a top–antitop pair	$pp \to t\bar{t}H$	y_t	direct top-Yukawa access but complex final states
tH	Higgs production with a single top quark	$pp \to tHj$	interference involving top and weak couplings	rarer and more advanced; sensitive to coupling signs and interference

Table 3: Main Standard Model Higgs-production modes at a proton–proton collider, interpreted as MadGraph learning examples.

Associated production. A relatively direct set of tree-level examples is associated production with a weak boson. In MadGraph notation one may generate the following pedagogical processes:

```
# LHC associated VH production: pedagogical LO examples
import model sm
generate p p > z h
add process p p > w+ h
add process p p > w- h
```

```
output LHC_VH_examples
launch
```

These processes probe the Higgs coupling to weak bosons and provide useful examples because the associated weak boson can later be decayed into leptonic final states for cleaner event selections.

Top-associated Higgs production. The process $pp \to t\bar{t}H$ is conceptually important because it gives direct access to the top Yukawa coupling:

```
# LHC top-associated Higgs production
import model sm
generate p p > t t~ h
output LHC_ttH
launch
```

In a first exercise it is best to generate the inclusive hard process before adding top, weak-boson, or Higgs decay chains.

Electroweak VBF-like Higgs plus jets. A useful introductory VBF-like sample can be obtained by selecting electroweak Higgs plus two-jet production at tree level:

```
# LHC electroweak Hjj generation: VBF-like pedagogical example
import model sm
generate p p > h j j QCD=0
output LHC_VBF_Hjj_EW
launch
```

The restriction `QCD=0` is intended to remove QCD vertices and keep electroweak diagrams at this order. This is a generator-level selection, not a full experimental VBF definition. A realistic VBF analysis also imposes characteristic kinematic selections, for example on the invariant mass and rapidity separation of the two jets. Thus students should not identify the label Hjj by itself with an experimentally pure VBF signal.

Gluon fusion and the HEFT approximation. The most delicate introductory case is gluon fusion. In the full Standard Model, $gg \to H$ first appears through a loop, dominated by the top quark. It is therefore not an ordinary tree-level process in the Standard Model model file. For fast pedagogical exercises, one often uses an effective Higgs-gluon interaction model in which the heavy top loop has been replaced by a local effective vertex:

```
# LHC gluon fusion in an effective Higgs-gluon model.
# Check the exact HEFT model name available in the local MG5_aMC installation.
import model heft
generate p p > h
output LHC_ggF_HEFT
launch
```

This HEFT-style calculation is useful for learning the workflow and the dominant production mechanism, but it is not identical to the full finite-top-mass Standard Model prediction. It becomes less reliable in kinematic regions where the heavy-top approximation is not adequate, for example when hard scales are comparable to or larger than the top-quark mass.

Loop-induced generation in the full Standard Model is possible in suitable MadGraph5_aMC setups, but it is more advanced and version-dependent. A template sometimes used for loop-induced studies has the form

```
# Advanced loop-induced template: verify with the installed MG5_aMC version.
# Do not use as a guaranteed command without checking the manual.
import model loop_sm
# Example syntax may be version-dependent:
# generate g g > h [noborn=QCD]
```

For Module 5, the important lesson is conceptual: ggF is loop-induced in the full SM, while HEFT replaces that loop by an effective interaction for simplified calculations.

19.5 Higgs production at FCC-ee in MadGraph

In an electron–positron collider, the initial state is much cleaner than at a hadron collider. There are no proton PDFs in the simplest parton-level exercises, and the centre-of-mass energy is fixed by the two beam energies. This makes e^+e^- Higgs production an especially transparent environment for connecting the Higgs couplings to observable processes. The FCC-ee physics programme is built around precision studies at several centre-of-mass energies, including a Higgs-factory stage [5].

The main processes for the present module are

$$e^+e^- \to ZH, \tag{19.6}$$

$$e^+e^- \to \nu_e\bar{\nu}_e H, \tag{19.7}$$

$$e^+e^- \to e^+e^-H. \tag{19.8}$$

The first process is Higgs-strahlung and directly probes the HZZ interaction. The second is dominated by WW-fusion and probes HWW. The third contains ZZ-fusion topologies and gives an additional handle on the neutral-current Higgs coupling.

Higgs-strahlung. The cleanest beginner example is

```
# FCC-ee Higgs-strahlung example
import model sm
generate e+ e- > z h
output FCCee_ZH
launch
```

For a representative $\sqrt{s} = 240$ GeV exercise, the corresponding beam settings may be entered interactively at launch or edited in `run_card.dat`:

```
# Representative FCC-ee run-card settings for sqrt(s) = 240 GeV
```

```
set lpp1 0
set lpp2 0
set ebeam1 120
set ebeam2 120
```

For $\sqrt{s} = 250$ GeV, one would use `ebeam1 = 125` and `ebeam2 = 125`. These numbers are representative exercise choices; a realistic collider study must use the intended beam setup and account for machine effects.

WW fusion. The WW-fusion channel can be generated as

```
# FCC-ee WW-fusion Higgs production
import model sm
generate e+ e- > ve ve~ h
output FCCee_WWfusion_H
launch
```

This process becomes increasingly important at higher centre-of-mass energy and gives a direct handle on the HWW coupling.

ZZ fusion. The neutral-current fusion topology may be represented by

```
# FCC-ee ZZ-fusion Higgs production
import model sm
generate e+ e- > e+ e- h
output FCCee_ZZfusion_H
launch
```

For realistic lepton-collider predictions, initial-state radiation, beam-energy spread, beamstrahlung, and higher-order electroweak effects can matter. Those refinements are beyond the introductory purpose of this section.

19.6 Higgs production at the LHeC in MadGraph

A lepton–proton collider provides a different view of Higgs production. The lepton beam can radiate electroweak bosons, while the proton beam supplies quarks through PDFs. The characteristic Higgs-production topologies are charged-current and neutral-current vector-boson-fusion processes. In an LHeC-like setup, the charged-current process probes primarily HWW, while the neutral-current process probes primarily HZZ.

The charged-current channels are

$$e^- p \to \nu_e H j, \qquad e^+ p \to \bar{\nu}_e H j, \tag{19.9}$$

where the neutrino or antineutrino leads to missing energy. The neutral-current channels are

$$e^- p \to e^- H j, \qquad e^+ p \to e^+ H j, \tag{19.10}$$

where the scattered charged lepton remains visible. In both cases the final-state jet is associated with the struck quark line on the proton side.

Charged-current examples. For an electron beam, a simple charged-current Higgs-production example is

```
# LHeC-like charged-current Higgs production with an electron beam
import model sm
generate e- p > ve h j QCD=0
output LHeC_CC_Higgs_em
launch
```

The charge-conjugate positron-beam process is

```
# LHeC-like charged-current Higgs production with a positron beam
import model sm
generate e+ p > ve~ h j QCD=0
output LHeC_CC_Higgs_ep
launch
```

Neutral-current examples. The corresponding neutral-current processes are

```
# LHeC-like neutral-current Higgs production with an electron beam
import model sm
generate e- p > e- h j QCD=0
output LHeC_NC_Higgs_em
launch
```

and

```
# LHeC-like neutral-current Higgs production with a positron beam
import model sm
generate e+ p > e+ h j QCD=0
output LHeC_NC_Higgs_ep
launch
```

The restriction `QCD=0` selects electroweak production at tree level. As in the LHC VBF case, this is a useful generator-level definition for a learning exercise, not a substitute for a realistic signal-region analysis.

Representative asymmetric beam settings. For a representative LHeC-like setup with a 60 GeV lepton beam and a 7 TeV proton beam, one may use settings of the following form:

```
# Representative LHeC-like run-card settings.
# Here beam 1 is the lepton beam and beam 2 is the proton beam.
set lpp1 0
```

```
set lpp2 1
set ebeam1 60
set ebeam2 7000
set pdlabel lhapdf
set lhaid <chosen_PDF_ID>
```

Here `lpp1 = 0` means that beam 1 is treated as a non-hadronic beam, while `lpp2 = 1` means that beam 2 is a proton beam described by PDFs. The placeholder `<chosen_PDF_ID>` must be replaced by a PDF set installed through LHAPDF. The beam order in `run_card.dat` must match the process definition. If the proton is put in beam 1 and the lepton in beam 2, the settings must be swapped accordingly.

19.7 The basic MadGraph workflow

The practical workflow can be summarised as

$$\text{model} \longrightarrow \text{process} \longrightarrow \text{output directory} \longrightarrow \text{launch} \longrightarrow \text{cards} \longrightarrow \sigma_{\text{LO}} \longrightarrow \text{events}. \tag{19.11}$$

A typical session begins by starting `mg5_aMC`, importing a model, generating a process, and writing the process directory. The user then launches the run and checks the cards. The most important card for the present purpose is `run_card.dat`, which contains beam energies, PDF choices, cuts, scales, and the requested number of events. The `param_card.dat` stores model parameters such as masses, widths, and couplings.

File or output	Meaning
`run_card.dat`	beam setup, cuts, scale choices, PDF settings, and event number
`param_card.dat`	model parameters, including particle masses, widths, and couplings
`crossx.html`	human-readable cross-section summary for the generated run
`results.dat`	numerical run summary stored inside the process directory
`unweighted_events.lhe.gz`	compressed Les Houches Event file containing parton-level unweighted events

Table 4: Common MadGraph files and outputs encountered in a basic leading-order event-generation workflow.

For an introductory workflow, it is useful to generate inclusive production first, before adding decay chains. Decays can later be included either directly in the process syntax or with a dedicated decay tool such as MadSpin. If one proceeds further, parton-level events can be passed to a shower and hadronisation generator such as PYTHIA [8], and then to a fast detector-simulation tool such as Delphes [9]. Those steps belong naturally to a later collider-phenomenology or data-analysis module.

19.8 From cross sections to event samples

A cross section and an event sample are related but not identical. A cross section is a rate-like quantity for a specified process and setup. An event sample is a list of representative final states distributed according to the matrix element, phase space, and cuts. The same generated process can therefore be used in two complementary ways: to report a leading-order cross section and to produce events for studying kinematic distributions.

Stage	Tool or output	Physical meaning
Matrix element	MadGraph	hard partonic scattering amplitude for the requested process
Integration	MadGraph	cross section for the chosen model, beams, cuts, scales, and parameters
Unweighted events	LHE file	parton-level event sample distributed according to the hard process
Decays	process syntax or MadSpin	resonance decay treatment, often within a narrow-width approximation
Shower and hadronisation	PYTHIA or similar	QCD/QED radiation, parton shower, hadron formation, and unstable-particle decays
Detector simulation	Delphes or experiment-specific simulation	approximate detector response and reconstructed objects
Analysis	ROOT, Python, Rivet, or similar	selections, histograms, fiducial regions, and observables

Table 5: Conceptual separation between cross-section calculation, parton-level event generation, showering, detector simulation, and analysis.

It is also useful to distinguish total and fiducial cross sections. A total cross section refers to the rate within the generator-level definition of the process, possibly without analysis cuts. A fiducial cross section refers to the rate after a specified set of kinematic and object-level selections. In realistic analyses, fiducial definitions are tied to detector acceptance, reconstruction, and background rejection. In this module, cuts should be introduced only as simple illustrative tools.

19.9 Validation and interpretation

Before interpreting a MadGraph number, students should check the setup carefully. The most important questions are:

- Which model has been imported: the full SM, an effective HEFT model, or another model?
- Is the process tree-level, loop-induced, or generated with an effective interaction?
- Are the particle names and charge-conjugate states correct?
- Are the beam energies and beam order correct?
- For proton beams, is the intended PDF set being used?
- Are the factorisation and renormalisation scales appropriate for the exercise?
- Are cuts present, and if so, do they define an inclusive or fiducial result?
- Are masses and widths in `param_card.dat` consistent with the intended Standard Model setup?
- Are decays included, or is the generated process inclusive in the unstable particles?

Important caution. Basic MadGraph runs are not official precision predictions for Higgs production. This is especially important for LHC Higgs physics, where higher-order QCD and electroweak corrections are essential and where gluon fusion is loop-induced. For precision values of Higgs production cross sections and branching ratios, one should use dedicated calculations and official summaries such as the LHC Higgs Cross Section Working Group results [4]. The purpose of the present section is more modest: to identify the mechanisms, learn the workflow, connect Higgs couplings to event topologies, and understand how cross sections and parton-level events are generated.

Take-home message
MadGraph is an excellent pedagogical tool for connecting the Higgs interactions in the Lagrangian to collider processes. Its output becomes physically meaningful only after the model, beam setup, PDFs, cuts, scales, decays, and perturbative accuracy have been specified.

19.10 Guided exercises

Guided checks

1. Generate `e+ e- > z h` at $\sqrt{s} = 240$, 250, and 365 GeV. Learning target: see how the Higgs-strahlung cross section depends on centre-of-mass energy.

2. Generate `e+ e- > ve ve~ h` and compare its qualitative energy dependence with `e+ e- > z h`. Learning target: understand why fusion processes become more important at higher energy.

3. Generate `p p > z h` and `p p > t t~ h` at an LHC-like energy. Learning target: identify which Higgs coupling each process tests most directly.

4. Generate an electroweak `p p > h j j QCD=0` sample and inspect the effect of simple VBF-like cuts on the two jets. Learning target: distinguish generator-level electroweak Hjj production from a full experimental VBF selection.

5. Generate `e- p > ve h j QCD=0` and `e- p > e- h j QCD=0` in an LHeC-like setup. Learning target: compare charged-current and neutral-current Higgs-production topologies.

6. Open one generated LHE file and identify the incoming beams, final-state particles, and event weights. Learning target: connect the generator output to the physical event interpretation.

This practical bridge completes the module's path from the electroweak Higgs mechanism to observable Higgs production. The Higgs couplings derived from symmetry breaking are not merely formal consequences of the Lagrangian; they determine the production channels, event topologies, and collider strategies through which the Higgs boson is studied.

20 What the Higgs sector teaches us about the Standard Model

After the bounded bridge to collider visibility, it is useful to return explicitly to the structural lessons of the module. The main point is not the existence of one more observed particle, but what the Higgs sector teaches us about the internal logic of the Standard Model.

20.1 Mass generation without abandoning gauge symmetry

The most important lesson of the Higgs sector is that mass generation need not come at the price of abandoning gauge symmetry. The Standard Model does not simply choose between symmetry and mass. Instead, it preserves symmetry in the Lagrangian and allows the vacuum to realize it nontrivially. That is the conceptual content of spontaneous symmetry breaking in the electroweak theory.

20.2 The Higgs sector as a structural requirement, not an optional add-on

It is tempting, especially after the dramatic history of the 2012 discovery, to think of the Higgs sector as a particle-discovery topic attached to the Standard Model from outside. That is the wrong conceptual order. The electroweak theory already demanded a mechanism of this type once one asked for a chiral gauge theory with massive weak bosons, massive fermions, and good high-energy behaviour. The observed Higgs boson matters so much because it is the physical manifestation of that structural solution.

20.3 Hidden symmetry versus destroyed symmetry

The phrase "broken symmetry" can be misleading if read too literally. The electroweak gauge symmetry is not destroyed in the sense that the Lagrangian forgets it. Rather, it is hidden by the chosen vacuum. The theory retains the symmetry in its underlying formulation, but the particle spectrum around the vacuum does not display the full symmetry manifestly. This distinction between hidden and explicitly destroyed symmetry is one of the most important conceptual takeaways of the module.

20.4 What the Higgs mechanism explains

Within the Standard Model, the Higgs mechanism explains why the $W^{\pm}$ and Z are massive while the photon is not, why charged fermion masses can arise in a gauge-invariant way, why Higgs couplings track particle masses, and why a neutral spin-zero scalar remains in the spectrum. It also supplies part of the high-energy consistency of the electroweak theory. These are substantial explanatory successes.

20.5 What remains unexplained: parameters, Yukawas, and flavour hierarchy

At the same time, the Higgs sector does not explain everything. It does not predict the values of the Yukawa couplings. It does not explain the pattern of fermion masses and mixings. It does not by itself explain why the Higgs potential has the numerical parameters it does. These limitations are not failures of the mechanism; they are signs that the Standard Model, despite its remarkable success, may not be the final theory.

Remark 20.1: What the Higgs mechanism does *not* explain

The Higgs mechanism explains *how* gauge-boson and charged-fermion masses arise consistently within the Standard Model. It does not explain *why* the Yukawa couplings have their observed numerical values, nor why the fermion mass hierarchy is so large. Keeping this distinction clear is essential for understanding both the power and the limitations of the Standard Model.

21 A bounded outlook: what Module 5 will not attempt

The preceding sections already contain the core conceptual content of Module 5. The present section merely marks nearby topics that are important, but that should remain outside the present note's main development.

21.1 Precision Higgs coupling fits

Modern Higgs physics includes increasingly precise measurements of production rates, decay rates, differential distributions, and coupling combinations. Such precision fits are extremely important, but they require a machinery of cross sections, loop corrections, detector effects, and statistical inference that goes beyond the present module. Here it is enough to understand the structural origin of the couplings and the qualitative reason why they are experimentally testable.

21.2 Extended Higgs sectors

The minimal Standard Model contains one Higgs doublet. Many theories beyond the Standard Model enlarge the scalar sector: two-Higgs-doublet models, supersymmetric models, composite-Higgs models, and other extensions. These are important in modern research, but they belong to a later discussion. The purpose of the present module is to understand the minimal mechanism first and well.

21.3 Neutrino masses in the minimal versus extended picture

In the minimal Standard Model, neutrinos are massless because no right-handed neutrino field is introduced and no renormalisable Yukawa term can be written for them. Yet neutrino oscillation experiments show that neutrinos do have mass. This means that the minimal Standard Model is incomplete in the neutrino sector. Module 5 acknowledges this fact but does not attempt a full treatment; that belongs to a later discussion of Standard Model limitations and extensions.

21.4 Naturalness and hierarchy as later questions

The scalar sector raises deep questions about naturalness and the stability of the electroweak scale under quantum corrections. These questions motivate a great deal of beyond-the-Standard-Model model building. They are important, but conceptually they come after one has understood the Higgs mechanism itself. For that reason they are only signposted here and deferred to later modules or advanced courses.

22 Bridge to later modules

With the symmetry-breaking mechanism, masses, and basic Higgs interpretation now in place, the natural next step is to see where this material leads in the larger course structure.

22.1 From electroweak symmetry breaking to amplitudes and observables: Module 6

Once the spectrum and couplings of the electroweak theory are understood, the next natural step is to compute processes: decay widths, scattering amplitudes, cross sections, and experimentally measurable observables. Module 6 takes this step. The Higgs couplings derived here become ingredients in explicit production and decay calculations.

22.2 From Higgs couplings to open questions: Module 7

The Higgs sector also opens several conceptual questions. Why are the Yukawa couplings so hierarchical? Why is there one Higgs doublet and not more? How should one understand neutrino masses, the hierarchy problem, or the possible incompleteness of the scalar sector? These questions point directly toward the limitations and open problems of the Standard Model, which belong to Module 7.

22.3 From collider Higgs signatures to data analysis: Module 8

The collider-facing Higgs section of the present note was intentionally qualitative. A full analysis of Higgs data requires event selection, reconstruction, kinematic distributions, background estimation, and statistical interpretation. Those practical aspects connect naturally to a later laboratory or data-analysis module. In that sense Module 5 bridges not only to formal phenomenology but also to the logic of experimental analysis.

23 Final summary and conceptual map

23.1 The logical chain of Module 5

The logic of the module can be summarized in one chain. The electroweak Standard Model is a chiral gauge theory. Naive mass terms for gauge bosons and fermions are not gauge invariant. A scalar sector with a suitable potential allows spontaneous symmetry breaking. In the global case one gets massless Goldstone bosons; in the gauge case the Goldstone modes are absorbed and gauge bosons become massive. A single complex Higgs doublet with a nonzero vacuum expectation value breaks $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$ to $\mathrm{U}(1)_{\mathrm{em}}$. The $W^{\pm}$ and Z acquire masses, the photon remains massless, fermion masses arise through Yukawa couplings, and one physical scalar excitation remains: the Higgs boson. The later LHC discussion then provides a bounded qualitative bridge from this structure to experiment.

23.2 The Higgs sector is not just a particle-discovery story

The observed Higgs boson matters because it is tied to one of the central structural ideas of the Standard Model. Without the Higgs sector the electroweak theory would not describe the observed world in its known form. The discovery of the Higgs boson is therefore best understood not as the addition of a final decorative particle, but as experimental access to the mechanism that reconciles gauge symmetry with the masses of weak bosons and charged fermions.

23.3 What the student should now be able to see

After this module, the student should be able to explain the difference between symmetry of the Lagrangian and symmetry of the vacuum, the meaning of spontaneous symmetry breaking, why the global and local cases differ, why the photon remains massless while the $W^{\pm}$ and Z do not, why fermion masses require Yukawa couplings, and why the observed Higgs boson is the physical scalar excitation of the broken electroweak vacuum. The student should also be able to place the main Higgs production channels at the LHC into the broader logic of the Standard Model.

Take-home message

The electroweak Standard Model cannot simply assign masses by hand. The Higgs mechanism solves this problem by keeping the Lagrangian gauge invariant while allowing the vacuum to choose a symmetry-breaking direction. As a result, the $W^{\pm}$ and Z become massive, the photon remains massless, charged fermion masses arise from Yukawa couplings, and one physical scalar — the Higgs boson — remains in the spectrum. The Higgs boson observed at the LHC is therefore the experimental manifestation of a structural ingredient of the Standard Model, not an isolated add-on.

A Conventions and notation summary

Quantity	Convention
Metric	$\eta_{\mu\nu} = \mathrm{diag}(1, -1, -1, -1)$
Natural units	$\hbar = c = 1$
Electroweak gauge group	$\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$
Electric charge	$Q = T_3 + Y$
Higgs doublet	$\Phi = (\phi^+, \phi^0)^T$
Weak mixing angle	θ_W with $\tan\theta_W = g'/g$
Vacuum expectation value	$\langle\Phi\rangle = \frac{1}{\sqrt{2}}(0, v)^T$

Convention note: this note uses $Q = T_3 + Y$. Some texts instead write $Q = T_3 + Y/2$; this corresponds to a hypercharge convention differing by a factor of two.

B Electroweak gauge and Higgs-sector formula summary

The core formulas of the module are collected here for quick reference:

$$D_\mu = \partial_\mu - \mathrm{i}g\frac{\tau^a}{2}W^a_\mu - \mathrm{i}g'YB_\mu, \tag{B.1}$$
$$V(\Phi) = \mu^2\Phi^\dagger\Phi + \lambda(\Phi^\dagger\Phi)^2, \tag{B.2}$$
$$\langle\Phi\rangle = \frac{1}{\sqrt{2}}\begin{pmatrix}0\\ v\end{pmatrix}, \tag{B.3}$$
$$v^2 = -\frac{\mu^2}{\lambda}, \tag{B.4}$$
$$m_W = \frac{gv}{2}, \tag{B.5}$$
$$m_Z = \frac{v}{2}\sqrt{g^2 + g'^2}, \tag{B.6}$$
$$m_h^2 = 2\lambda v^2. \tag{B.7}$$

C Photon–Z mixing and charge-assignment summary

The neutral electroweak mass matrix is

$$\frac{v^2}{8}\begin{pmatrix} g^2 & -gg' \\ -gg' & g'^2 \end{pmatrix}. \tag{C.1}$$

It is diagonalised by the Weinberg rotation,

$$A_\mu = \sin\theta_W\, W^3_\mu + \cos\theta_W\, B_\mu, \qquad Z_\mu = \cos\theta_W\, W^3_\mu - \sin\theta_W\, B_\mu. \tag{C.2}$$

The photon remains massless because the unbroken generator is

$$Q = T_3 + Y. \tag{C.3}$$

The electric coupling is related to the electroweak couplings by

$$e = g\sin\theta_W = g'\cos\theta_W. \tag{C.4}$$

D Yukawa and fermion-mass summary

Representative Yukawa terms for one generation are

$$\mathcal{L}_Y = -y_e\,\bar{L}_L\Phi e_R - y_d\,\bar{Q}_L\Phi d_R - y_u\,\bar{Q}_L\widetilde{\Phi}u_R + \text{h.c.}, \tag{D.1}$$
$$\widetilde{\Phi} = \mathrm{i}\tau^2\Phi^*. \tag{D.2}$$

After electroweak symmetry breaking,

$$m_f = \frac{y_f v}{\sqrt{2}}, \qquad \mathcal{L}_{hff} = -\frac{m_f}{v}h\bar{f}f. \tag{D.3}$$

These equations summarize the Standard Model origin of charged-fermion masses and Higgs–fermion couplings.

E Higgs production and decay channel summary

Category	Channel	Main pedagogical point
Production	ggF	Dominant rate; gluon luminosity plus top-loop mechanism
Production	VBF	Cleaner topology; probes Higgs coupling to weak bosons
Production	VH	Tagged environment using associated W/Z boson
Production	ttH	Direct sensitivity to the top Yukawa coupling
Decay / search	$H \to \gamma\gamma$	Rare but very clean discovery channel
Decay / search	$H \to ZZ^* \to 4\ell$	Extremely clean channel with precise mass reconstruction
Decay / search	$H \to WW^*$	Important bosonic decay mode
Decay / search	$H \to \tau^+\tau^-$	Direct probe of a lepton Yukawa coupling
Decay / search	$H \to b\bar{b}$	Largest branching channel near 125 GeV but experimentally challenging

F Guided-check summary

Guided checks

- Can you explain why a direct gauge-boson mass term is not gauge invariant?
- Can you distinguish clearly between symmetry of the Lagrangian and symmetry of the vacuum?
- Can you describe what happens to the four real degrees of freedom of the Higgs doublet after electroweak symmetry breaking?
- Can you explain why the photon remains massless but the Z does not?
- Can you explain why fermion masses require Yukawa couplings in the electroweak Standard Model?
- Can you name the main Higgs production modes at the LHC and state qualitatively why gluon fusion dominates?

References

[1] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli and M. Zaro, “The automated computation of tree-level and next-to-leading order

differential cross sections, and their matching to parton shower simulations," JHEP **07**, 079 (2014) doi:10.1007/JHEP07(2014)079 [arXiv:1405.0301 [hep-ph]].

[2] S. Dittmaier *et al.* [LHC Higgs Cross Section Working Group], "Handbook of LHC Higgs Cross Sections: 1. Inclusive Observables," CERN Yellow Report **CERN-2011-002** (2011) doi:10.5170/CERN-2011-002 [arXiv:1101.0593 [hep-ph]].

[3] S. Heinemeyer *et al.* [LHC Higgs Cross Section Working Group], "Handbook of LHC Higgs Cross Sections: 3. Higgs Properties," CERN Yellow Report **CERN-2013-004** (2013) doi:10.5170/CERN-2013-004 [arXiv:1307.1347 [hep-ph]].

[4] D. de Florian *et al.* [LHC Higgs Cross Section Working Group], "Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector," CERN Yellow Rep. Monogr. **2**, 1–869 (2017) doi:10.23731/CYRM-2017-002 [arXiv:1610.07922 [hep-ph]].

[5] M. Benedikt, F. Zimmermann *et al.*, "Future Circular Collider Feasibility Study Report: Volume 1, Physics, Experiments, Detectors," CERN-FCC-PHYS-2025-0002 [arXiv:2505.00272 [hep-ex]].

[6] J. L. Abelleira Fernandez *et al.* [LHeC Study Group], "A Large Hadron Electron Collider at CERN: Report on the Physics and Design Concepts for Machine and Detector," J. Phys. G **39**, 075001 (2012) doi:10.1088/0954-3899/39/7/075001 [arXiv:1206.2913 [physics.acc-ph]].

[7] A. Buckley, J. Ferrando, S. Lloyd, K. Nordström, B. Page, M. Rüfenacht, M. Schönherr and G. Watt, "LHAPDF6: parton density access in the LHC precision era," Eur. Phys. J. C **75**, 132 (2015) doi:10.1140/epjc/s10052-015-3318-8 [arXiv:1412.7420 [hep-ph]].

[8] C. Bierlich *et al.*, "A comprehensive guide to the physics and usage of PYTHIA 8.3," SciPost Phys. Codebases **8**, 1 (2022) doi:10.21468/SciPostPhysCodeb.8 [arXiv:2203.11601 [hep-ph]].

[9] J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaître, A. Mertens and M. Selvaggi [DELPHES 3], "DELPHES 3: A modular framework for fast simulation of a generic collider experiment," JHEP **02**, 057 (2014) doi:10.1007/JHEP02(2014)057 [arXiv:1307.6346 [hep-ex]].