AGH University of Krakow

Faculty of Physics and Applied Computer Science

MODULE 1

MATHEMATICAL TOOLS AND FORMAL LANGUAGE

Supporting Lecture Notes for *The Standard Model*

The Guideline

The aim of this module is not to accumulate abstract definitions for their own sake. It is to identify the minimal mathematical language needed to understand why the Standard Model has the form it does. States, operators, symmetries, groups, generators, representations, and invariants are not decorative tools. They are the grammar in which modern particle physics is written.

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

1 Introduction and motivation

Module 1 provides the formal language on which the rest of the course is built. Before one can discuss relativistic wave equations, Lorentz symmetry, gauge fields, or the structure of the Standard Model Lagrangian, one must first understand the mathematical objects that make such constructions possible. For this reason, the purpose of the present module is to introduce the basic framework of states, observables, linear operators, eigenvalue problems, commutators, symmetries, invariants, continuous groups, generators, Lie algebras, and representations.

The motivation for beginning in this way is simple but important. The Standard Model is not best understood as a catalogue of particles and interactions to be memorised. It is a highly constrained relativistic quantum field theory whose structure is controlled by symmetry. The allowed fields, the possible interaction terms, the conserved quantities, and eventually the appearance of gauge fields are all tied to precise mathematical principles. In this sense, the formal language developed in Module 1 is not auxiliary background. It is part of the internal logic of the theory itself.

The module begins with the quantum-mechanical description of physical systems in terms of Hilbert space, state vectors, and operators. This is the natural starting point because quantum field theory extends, rather than replaces, the operator language of quantum mechanics. We then move to operator algebra and commutators, where one first sees that physical information is encoded not only in observables themselves but also in the relations among them. From there the discussion turns to symmetry and invariance, which provide the bridge from formal structure to physical law.

A central theme of the module is that symmetry is not merely a convenient property of a theory that has already been written down. In modern particle physics, symmetry is a constructive principle. Continuous symmetries lead to conserved quantities; groups and Lie algebras organise the transformations systematically; representations tell us how physical fields transform; and invariant combinations determine which terms may appear in a Lagrangian. Once this logic is understood, the transition from global symmetry to local symmetry naturally motivates the introduction of gauge fields and prepares the conceptual path toward gauge theory.

The aim of the module is therefore twofold. On the one hand, it develops a self-contained introduction to the formal tools that recur throughout particle physics. On the other hand, it prepares the ground for the next stages of the course, where relativistic symmetry, spinors, and gauge interactions will be studied in detail. By the end of Module 1, the reader should be able to see the Standard Model not as a disconnected collection of formulas, but as a theory whose architecture follows from a small set of powerful mathematical ideas.

The Guideline

Central pedagogical question. What is the minimal mathematical language needed to construct the Standard Model in a logically coherent way? The answer will unfold through the chain

$$\begin{aligned}\text{states} &\to \text{operators} \to \text{commutators} \to \text{symmetry}\\ &\to \text{groups and generators} \to \text{representations}\\ &\to \text{invariants} \to \text{local gauge structure.}\end{aligned}$$

1.1 Why formal language comes first

Students often meet quarks, leptons, gauge bosons, and the Higgs boson before they meet the mathematical logic that organises them. That first exposure is useful, but it can easily leave the impression that particle physics is an inventory. The Standard Model becomes intelligible only when one asks structural questions:

- What is a physical state?
- What counts as an observable?
- What does it mean for operators to be compatible or incompatible?
- What is a symmetry and what is an invariant?
- Why do continuous symmetries lead to conserved quantities?
- How are transformations organised mathematically?
- Why do particle multiplets appear at all?
- Why does local symmetry force gauge fields?

These are not side questions. They are exactly the questions that make the Standard Model look like a coherent theory rather than a collection of phenomenological facts.

1.2 Roadmap of these notes

We begin with the quantum-mechanical language of states and observables. We then examine commutators and explain why algebra itself contains physical content. Next we introduce symmetry and invariance, followed by a light but meaningful account of Noether's theorem. We then develop the language of groups, generators, and Lie algebras. After that, we discuss representations and multiplets and show why they matter for particle assignments. We then turn to invariants and the logic of Lagrangian building. Finally, we explain why the transition from global to local symmetry leads naturally to covariant derivatives and gauge fields, providing a first roadmap to the Standard Model and a bridge to Modules 2 and 3.

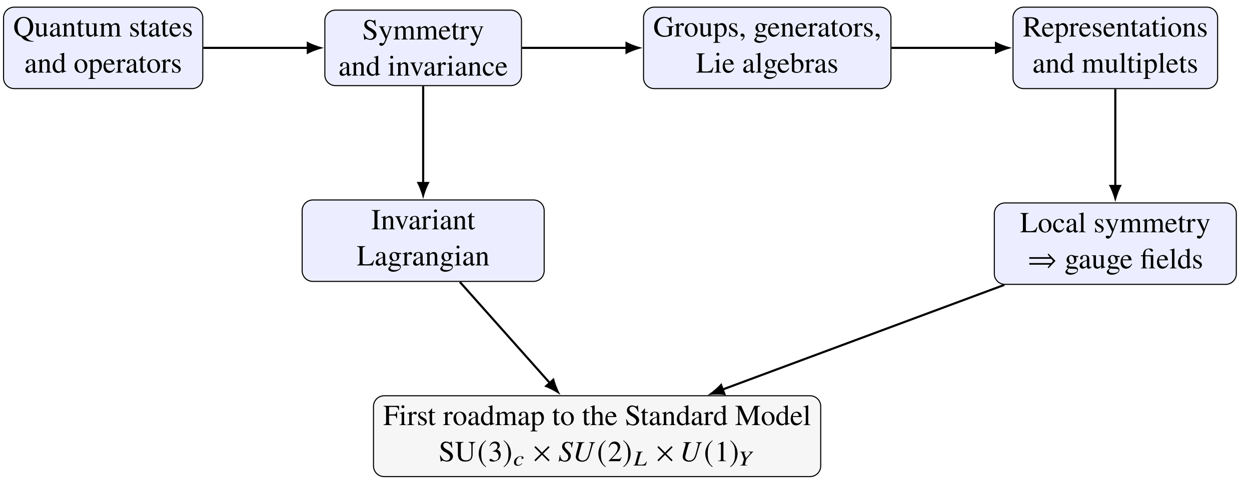

Figure 1: Roadmap of Module 1. The module moves from the operator language of quantum theory to the symmetry-based logic that points toward the Standard Model gauge structure.

2 From quantum mechanics to the language of modern particle physics

2.1 Why quantum mechanics is the right starting point

Ordinary quantum mechanics is the first successful mathematical language in which a physical system is described by a state vector, measurable quantities are represented by operators, and probabilities are obtained from inner products and expectation values. Even when we later pass to quantum fields, the basic logic remains recognisable. Quantum field theory does not discard the operator language of quantum mechanics; it extends it so that fields themselves become operator-valued objects and particle number need no longer be fixed.

This point is worth pausing over because students sometimes think quantum mechanics and quantum field theory are almost disjoint subjects. That is not the right conceptual picture. Quantum field theory keeps the basic logic of quantum mechanics:

- superposition of states remains true,
- observables are still operator-valued,
- symmetry is still implemented by operators,
- conserved quantities are still associated with invariance,
- algebraic relations among operators still matter.

What changes is the domain of application. Particle physics is relativistic, particles can be created and destroyed, and interactions are most naturally expressed in terms of fields living over spacetime. But before we study relativistic fields, we need the common formal language that survives the transition.

Remark 2.1: What Module 1 does not yet do

Module 1 does not derive the full Standard Model and does not assume a full prior course in quantum field theory. Its task is more basic: to explain why later relativistic and gauge-theoretic constructions must be written in a precise mathematical language and why that language is already visible in simpler quantum systems.

2.2 Why particle physics needs more than non-relativistic quantum mechanics

There are at least three reasons why particle physics cannot stop at the level of non-relativistic quantum mechanics.

1. **Relativity.** The theory must respect spacetime symmetry. This will become the central topic of Module 2.
2. **Variable particle number.** Relativistic processes can create and annihilate particles. A theory with permanently fixed particle number is too restrictive.
3. **Local interactions.** The natural way to express interactions compatible with locality is through fields defined at spacetime points.

These three facts explain why quantum field theory is the natural language of the Standard Model. Yet they do not change the foundational role of states, operators, and symmetries. Rather, they show why those ideas must be generalised carefully.

2.3 The first look ahead

The first three modules of the course fit together naturally.

Module	Conceptual role
Module 1	Gives the mathematical grammar: states, operators, commutators, symmetry, groups, generators, representations, invariants, and the global-to-local transition.
Module 2	Adds the relativistic spacetime framework: Lorentz symmetry, four-vectors, spinors, Dirac matrices, and the Dirac equation.
Module 3	Uses the language of Modules 1 and 2 to construct gauge theories and then the gauge structure of the Standard Model.

The whole purpose of Module 1 is therefore to make Modules 2 and 3 feel like natural continuations rather than abrupt changes of topic.

3 State space and observables

3.1 Hilbert space as the arena of states

In quantum theory, the state of a physical system is represented by a vector in a complex Hilbert space. Informally, a Hilbert space is a complex vector space equipped with an inner product and complete enough to support the limiting procedures needed in quantum theory.

Definition 3.1: Hilbert space and state vector

A *Hilbert space* $\mathcal{H}$ is a complex vector space with inner product $\langle \phi \mid \psi \rangle$ such that the induced norm

$$\|\psi\| = \sqrt{\langle \psi \mid \psi \rangle}$$

makes $\mathcal{H}$ complete. A physical *state* is represented by a ray in $\mathcal{H}$, typically described by a normalised vector $|\psi\rangle$ with $\langle \psi \mid \psi \rangle = 1$.

The distinction between a vector and a ray matters. The vectors $|\psi\rangle$ and $e^{i\alpha}|\psi\rangle$ represent the same physical state because the overall phase is not observable. Already here quantum theory teaches us that not every mathematical feature of the description is physically meaningful.

Example 3.1: A finite-dimensional example: spin one-half

The Hilbert space of a spin one-half system is $\mathbb{C}^2$. A convenient basis is

$$|\uparrow\rangle = \begin{pmatrix}1\\0\end{pmatrix}, \qquad |\downarrow\rangle = \begin{pmatrix}0\\1\end{pmatrix}.$$

A general normalised state is

$$|\psi\rangle = a|\uparrow\rangle + b|\downarrow\rangle, \qquad |a|^2 + |b|^2 = 1.$$

This simple system already contains superposition, measurement probabilities, operator action, and non-commuting observables.

3.2 Bras, kets, and inner products

Dirac notation is compact because it displays the algebraic structure of the theory clearly. To each state vector $|\psi\rangle$ there corresponds a dual vector $\langle\psi|$, and the inner product is written as $\langle\phi \mid \psi\rangle$. The physical role of the inner product is twofold. It allows us to normalise states, and it governs probability amplitudes. If $|a\rangle$ is an eigenstate of an observable, then $|\langle a \mid \psi\rangle|^2$ is the probability that a measurement in the state $|\psi\rangle$ yields the outcome associated with $|a\rangle$.

The algebra becomes especially transparent once one writes the orthonormality and completeness relations explicitly. For an orthonormal basis $\{|n\rangle\}$ one has

$$\langle m \mid n\rangle = \delta_{mn}, \qquad \sum_n |n\rangle\langle n| = \mathbf{1},$$

where $\mathbf{1}$ is the identity operator on the Hilbert space. The second relation is not an abstract ornament: it says that the basis is complete, so every state can be reconstructed from its components in that basis. Acting with the identity on a state gives

$$|\psi\rangle = \mathbf{1}|\psi\rangle = \sum_n |n\rangle\langle n \mid \psi\rangle = \sum_n c_n|n\rangle, \qquad c_n \equiv \langle n \mid \psi\rangle.$$

The coefficients c_n are probability amplitudes. If the basis is the eigenbasis of an observable, then $|c_n|^2$ gives the probability of measuring the corresponding eigenvalue.

In particle physics this notation becomes especially useful because we do not want a formalism tied to one specific coordinate representation. Transformation properties matter more than a particular wavefunction written in position space. Dirac notation lets us talk directly about states, amplitudes, and operators without committing too early to position-space wavefunctions or matrix components.

3.3 States, bases, and superposition in practice

The abstract language of Hilbert space becomes less intimidating once one remembers that it is a natural generalisation of ordinary vector-space reasoning. A basis in the state space plays a role analogous to a basis in linear algebra. The difference is that basis vectors now correspond to physically meaningful alternatives. For example, in the spin one-half system, the basis $\{|\uparrow\rangle, |\downarrow\rangle\}$ gives a useful language for measurements of spin along a chosen axis. A general state is a superposition because quantum theory allows the system to be prepared in a coherent linear combination of basis states.

This superposition principle is one of the reasons why the operator language is indispensable. If a state can be a linear combination of simpler states, then the physical quantities that act on the state space must respect that linear structure. This is why observables are represented by linear operators rather than by arbitrary nonlinear rules. Concretely, if

$$|\psi\rangle = a|1\rangle + b|2\rangle,$$

and if $\hat{A}$ is linear, then

$$\hat{A}|\psi\rangle = a\hat{A}|1\rangle + b\hat{A}|2\rangle.$$

Linearity is therefore the mathematical reflection of the physical superposition principle.

A second important point is that the basis is a matter of description, not of physics. Different bases can be more or less convenient for different observables, but they describe the same state space. What is physical is not the particular column vector used in one basis, but the state itself and the probabilities extracted from it.

Example 3.2: A basis change already carries physics

Suppose a state is written in the z-spin basis as

$$|\psi\rangle = \frac{1}{\sqrt{2}}|\uparrow\rangle + \frac{1}{\sqrt{2}}|\downarrow\rangle.$$

In that basis it is a superposition. But in the basis of eigenstates of spin along the x axis, the same state may be written as

$$|+_x\rangle = \frac{1}{\sqrt{2}}(|\uparrow\rangle + |\downarrow\rangle), \qquad |-_x\rangle = \frac{1}{\sqrt{2}}(|\uparrow\rangle - |\downarrow\rangle),$$

so that $|\psi\rangle = |+_x\rangle$. The same physical state is a superposition in one basis and an eigenstate in another. This reminds us that a vector in Hilbert space is basis-independent, while its components depend on the basis chosen for the problem at hand.

3.4 Observables as linear Hermitian operators

An observable is represented by a linear operator acting on the Hilbert space of states. Linearity is essential because quantum states satisfy the superposition principle.

Definition 3.2: Observable and Hermitian operator

An *observable* is represented by a linear Hermitian operator $\hat{A}$ on $\mathcal{H}$. Hermiticity means

$$\hat{A}^\dagger = \hat{A},$$

which guarantees that its eigenvalues are real.

This is not just a formal convention. A measurement result must be a real number. Hermitian operators are the natural mathematical objects with that property. The basic reason can be seen directly from expectation values. If $\hat{A} = \hat{A}^\dagger$, then

$$\big(\langle\psi|\hat{A}|\psi\rangle\big)^* = \langle\psi|\hat{A}^\dagger|\psi\rangle = \langle\psi|\hat{A}|\psi\rangle,$$

so the expectation value is real. Likewise, if $\hat{A}|a\rangle = a|a\rangle$ and $\langle a \mid a\rangle \neq 0$, then

$$\langle a|\hat{A}|a\rangle = a\langle a \mid a\rangle = \langle a|\hat{A}^\dagger|a\rangle = a^*\langle a \mid a\rangle,$$

which implies $a = a^*$. Hence the eigenvalues of a Hermitian operator are real, as required for measurable quantities.

3.5 Eigenvalues, eigenvectors, and spectral meaning

The measurement content of an observable is encoded in the eigenvalue equation

$$\hat{A}|a\rangle = a|a\rangle.$$

If the system is in the eigenstate $|a\rangle$, then a measurement of A yields a with certainty. For a general state $|\psi\rangle$ one expands in the eigenbasis of $\hat{A}$, and the coefficients determine the probabilities of different outcomes.

In a finite-dimensional case this may be written schematically as

$$\hat{A} = \sum_n a_n |a_n\rangle\langle a_n|.$$

This is the spectral decomposition of the operator. It says that the action of the observable can be reconstructed from its eigenvalues and the projectors $|a_n\rangle\langle a_n|$. The point is not merely that operators act; they come with spectral data that encode what can actually be observed.

The same decomposition also clarifies measurement probabilities. If

$$|\psi\rangle = \sum_n c_n |a_n\rangle,$$

then a measurement of A returns one of the eigenvalues a_n with probability $|c_n|^2$. The formalism therefore links algebraic structure directly to experimental meaning.

3.6 Expectation values

The expectation value of an observable $\hat{A}$ in the state $|\psi\rangle$ is

$$\langle\hat{A}\rangle_\psi = \langle\psi|\hat{A}|\psi\rangle.$$

This number is the average result obtained from many measurements on identically prepared systems. It need not be one of the possible outcomes of a single measurement. If the state has the expansion $|\psi\rangle = \sum_n c_n|a_n\rangle$ in the eigenbasis of $\hat{A}$, then

$$\langle\hat{A}\rangle_\psi = \sum_n |c_n|^2 a_n.$$

Thus the expectation value is a probability-weighted average of the possible outcomes.

A related quantity is the variance,

$$\mathrm{Var}(A) = \langle \hat{A}^2 \rangle_\psi - \langle \hat{A} \rangle_\psi^2,$$

which measures how sharply or broadly the measurement outcomes are distributed. If the variance vanishes, the state is an eigenstate of the observable and the outcome is sharp. This is the cleanest way to distinguish a definite value from a mere average.

Remark 3.1: Expectation value versus eigenvalue

Students often confuse the expectation value with the value of the observable. They are not the same in general. A system may have a well-defined expectation value even when it is not in an eigenstate of the corresponding operator. For example, a spin state can have vanishing expectation value of S_z without the measurement outcome ever being exactly zero.

Expectation values will matter later in field theory when one speaks about the vacuum, order parameters, and spontaneous symmetry breaking. They therefore already deserve conceptual clarity here.

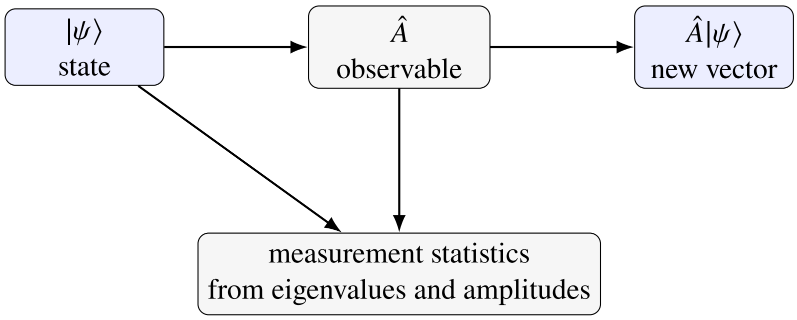

Figure 2: State-operator-measurement schematic. The state tells us what the system is, the operator tells us what can be measured, and the eigenstructure of the operator determines the possible outcomes.

3.7 Why this matters for the Standard Model

The Standard Model is a quantum field theory. In such a theory the fields are operator-valued objects acting on a Hilbert space of states. This means that the language of states and operators is not merely introductory background. It is part of the grammar of the theory itself. If one is not comfortable distinguishing the state, the operator, the eigenvalue, and the expectation value, later ideas such as currents, charges, vacuum structure, and field operators remain conceptually blurred.

4 Operators, algebra, and commutators

4.1 Operator products and non-commutativity

Once observables are represented by operators, physics is not encoded only in each operator separately. It is also encoded in how different operators relate to one another. The central object here is the commutator.

Definition 4.1: Commutator

For operators $\hat{A}$ and $\hat{B}$, the commutator is defined by

$$[\hat{A}, \hat{B}] = \hat{A}\hat{B} - \hat{B}\hat{A}.$$

If $[\hat{A}, \hat{B}] = 0$, the operators commute. If not, the order of the operations matters.

This algebraic quantity has direct physical content. Commuting observables can often be simultaneously diagonalised and are therefore mutually compatible. Non-commuting observables cannot in general be assigned sharp values simultaneously.

4.2 Compatible observables and simultaneous diagonalisation

If two Hermitian operators commute, one may typically choose a basis that diagonalises them simultaneously. This explains why states in quantum mechanics are often labelled by several quantum numbers at once. The existence of such simultaneous labels already reflects an underlying commuting algebra.

A useful proof sketch goes as follows. Suppose $\hat{A}$ and $\hat{B}$ commute and let $|a\rangle$ be an eigenvector of $\hat{A}$ with eigenvalue a. Then

$$\hat{A}\big(\hat{B}|a\rangle\big) = \hat{B}\hat{A}|a\rangle = a\,\hat{B}|a\rangle.$$

So $\hat{B}|a\rangle$ is again an eigenvector of $\hat{A}$ with the same eigenvalue. This means that $\hat{B}$ maps the eigenspace of $\hat{A}$ into itself. One may then diagonalise $\hat{B}$ inside each eigenspace of $\hat{A}$, producing a common eigenbasis whenever the usual spectral assumptions are satisfied. The physical conclusion is that commuting observables can be simultaneously sharp.

Example 4.1: Angular momentum labels

One labels states by $|\ell, m\rangle$ because

$$\left[L^2, L_z\right] = 0.$$

By contrast,

$$\left[L_x, L_y\right] = i\hbar L_z,$$

so different Cartesian components of angular momentum are not simultaneously sharp. The non-zero commutator tells us that a state of definite L_x cannot at the same time be a state of definite L_y.

4.3 The canonical example: position and momentum

The most famous commutator in quantum mechanics is

$$[\hat{x}, \hat{p}] = i\hbar.$$

This already contains important physics. It implies the uncertainty relation

$$\Delta x\,\Delta p \geq \frac{\hbar}{2},$$

but it also teaches a deeper lesson: the algebra is itself part of the content of the theory. Position and momentum are not independent classical labels hiding inside the quantum formalism. Their non-commuting structure is fundamental.

The canonical commutator can also be understood as the quantum version of translations. In the position representation one has $\hat{p} = -\mathrm{i}\hbar\partial_x$, and acting on a test function $f(x)$ gives

$$(\hat{x}\hat{p} - \hat{p}\hat{x})f(x) = x(-\mathrm{i}\hbar f') - \big(-\mathrm{i}\hbar\partial_x(xf)\big) = -\mathrm{i}\hbar x f' + \mathrm{i}\hbar(f + xf') = \mathrm{i}\hbar f(x).$$

Since this holds for arbitrary $f(x)$, one obtains $[\hat{x}, \hat{p}] = \mathrm{i}\hbar$. The point is worth noticing: the commutator is not an extra axiom glued onto the theory after the fact. It is tied to the very way translations act on wavefunctions.

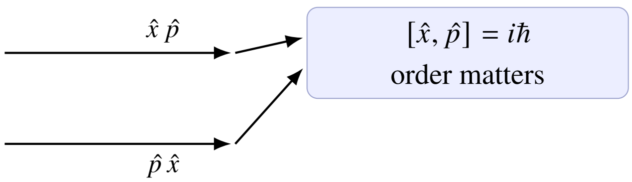

Figure 3: Schematic illustration of a non-zero commutator. The two operator orderings $\hat{x}\,\hat{p}$ and $\hat{p}\,\hat{x}$ are not the same, and their difference is encoded in the commutator $[\hat{x}, \hat{p}] = i\hbar$.

4.4 Commutators as a first bridge to symmetry generators

At this stage a student may ask why this operator algebra matters for a Standard Model course. The answer is that the same logic reappears in the mathematics of symmetry. Rotations, internal transformations, and gauge symmetries are all associated with generators, and the relations among these generators are encoded in commutators. In other words, the commutator language of elementary quantum mechanics becomes the Lie-algebra language of symmetry in field theory.

For example, if G is the generator of an infinitesimal transformation, then the transformed operator is obtained from

$$\hat{O}' = U\hat{O}U^{-1}, \qquad U = 1 + \mathrm{i}\epsilon G + O(\epsilon^2).$$

Expanding to first order gives

$$\delta\hat{O} \equiv \hat{O}' - \hat{O} = \mathrm{i}\epsilon\left[G, \hat{O}\right] + O(\epsilon^2).$$

So the commutator with the generator tells us how the observable changes under the symmetry. This is exactly the kind of structure that later becomes central in Noether charges, Lie algebras, and gauge theory.

Take-home message

A theory is not determined only by its list of observables. It is also determined by the algebra those observables satisfy. In later Standard Model language, much of the structure of the theory will be encoded not only in fields and particles, but in the commutators of the symmetry generators.

5 Symmetry and invariance

5.1 What is a symmetry?

The word symmetry is used so often that its meaning can seem almost obvious. But in theoretical physics it must be stated carefully. A symmetry is a transformation under which the physical content of the system is unchanged. The transformation may act on coordinates, on fields, on internal labels, or on the state space. What matters is that observable physics remains the same.

Definition 5.1: Symmetry and invariant quantity

A *symmetry* is a transformation under which the physical description of a system remains unchanged. A quantity is called *invariant* if its value is unchanged by that transformation.

This already contains two ideas that students sometimes blur together: the transformation and the invariant. A symmetry tells us what operation is performed. An invariant tells us what survives that operation unchanged.

5.2 Active and passive viewpoints

A symmetry can often be interpreted actively or passively. In an active viewpoint one transforms the system itself. In a passive viewpoint one changes the coordinate description. In many physical contexts the two viewpoints are mathematically equivalent and philosophically complementary. For Module 1 the key lesson is that physical predictions do not depend on arbitrary descriptive choices.

5.3 Unitary implementation of symmetry

In quantum mechanics a symmetry transformation is often represented by a unitary operator U acting on the state space. The state changes as

$$|\psi\rangle \to U|\psi\rangle,$$

while an observable transforms as

$$\hat{A} \to U\hat{A}U^{-1}.$$

The transformation is called a symmetry if the measurable content of the theory is unchanged. This language is particularly useful because it links the abstract concept of symmetry directly to operator algebra.

Why unitary? The short answer is that transition probabilities must be preserved. If two states are transformed as $|\phi\rangle \to U|\phi\rangle$ and $|\psi\rangle \to U|\psi\rangle$, then the probability amplitude between them should not change in absolute value. This requires

$$\langle \phi \mid \psi \rangle \longrightarrow \langle \phi|U^\dagger U|\psi\rangle = \langle \phi \mid \psi \rangle \qquad \text{for all } |\phi\rangle, |\psi\rangle,$$

which implies

$$U^\dagger U = \mathbf{1}.$$

In quantum mechanics, probability-preserving symmetries are represented by unitary or antiunitary operators (Wigner's theorem). For the continuous symmetries connected to the identity that are the main focus here, the implementation is unitary. The statement is physically important: the symmetry operation may change the description of the state, but it cannot change the probability interpretation of the theory.

For continuous symmetries one writes

$$U(\alpha) = \exp(\mathrm{i}\alpha T),$$

where T is the generator. Expanding for small α gives

$$U(\alpha) = 1 + \mathrm{i}\alpha T + O(\alpha^2).$$

If $U(\alpha)$ is unitary, then

$$U(\alpha)^\dagger U(\alpha) = \mathbf{1} \quad \Rightarrow \quad (1 - \mathrm{i}\alpha T^\dagger)(1 + \mathrm{i}\alpha T) = \mathbf{1} + O(\alpha^2),$$

which at first order implies

$$T^\dagger = T.$$

So the generator of a continuous unitary symmetry is Hermitian. This is the same pattern encountered earlier for observables, and it is one of the reasons generators and conserved charges behave so much like measurable quantities.

One can also see why the exponential form is natural. Performing N infinitesimal transformations of size α/N gives

$$\left(1 + \mathrm{i}\frac{\alpha}{N}T\right)^N \xrightarrow[N\to\infty]{} \exp(\mathrm{i}\alpha T).$$

A finite continuous symmetry is therefore the accumulated effect of many infinitesimal ones. Later, when fields rather than simple quantum states become the main dynamical variables, exactly the same logic survives.

5.4 Discrete and continuous symmetries

Some symmetries occur in isolated steps, others depend smoothly on continuous parameters.

- **Discrete symmetries:** parity P, charge conjugation C, time reversal T.
- **Continuous symmetries:** rotations, translations, phase symmetries such as $\psi \to e^{\mathrm{i}\alpha}\psi$.

Continuous symmetries are especially important for the Standard Model because they lead to generators, conserved currents, and Lie algebras. Discrete symmetries remain physically significant, but they play a different structural role.

5.5 Spacetime versus internal symmetry

A second essential distinction is between symmetries acting on spacetime and symmetries acting in an internal space.

- **Spacetime (external) symmetries:** translations, rotations, Lorentz transformations.
- **Internal symmetries:** isospin, weak isospin, hypercharge, color.

The Standard Model combines both. Already at this early stage students should become comfortable with the fact that spacetime symmetry and internal gauge symmetry play conceptually different roles.

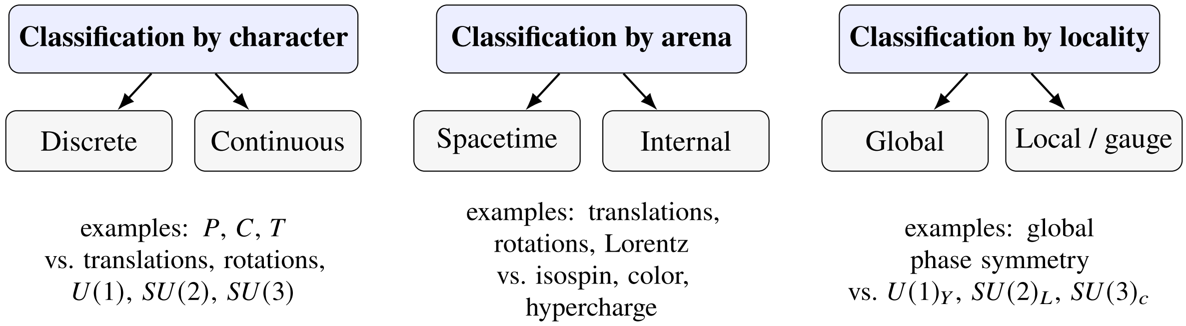

Figure 4: Different ways of classifying symmetries in modern particle physics. The labels discrete/continuous, spacetime/internal, and global/local refer to distinct classification axes, not to a single parent–child hierarchy. In the Standard Model, the central internal gauge structure is based on $SU(3)_c \times SU(2)_L \times U(1)_Y$.

5.6 Why symmetry is not decoration

In elementary mechanics symmetry is often introduced as a way to simplify a problem that is already solved. In modern particle physics it plays a deeper role. It is not only a property of a finished theory. It is also a principle used to build the theory. Once the required symmetry is specified, many possible terms in the Lagrangian are immediately forbidden. This is one reason the Standard Model is so constrained.

Remark 5.1: A change of viewpoint

A major conceptual step toward the Standard Model is the shift from *symmetry as a property of a finished system* to *symmetry as a principle used to construct the theory itself.*

6 Symmetry and conservation laws

6.1 The basic statement of Noether's theorem

One of the deepest results in theoretical physics is Noether's theorem. For the level of Module 1, the essential statement is simple: every continuous global symmetry of the action is associated with a conserved quantity. This is the bridge between symmetry and dynamics.

In a field theory with action

$$S = \int \mathrm{d}^4x \, \mathcal{L}(\phi, \partial_\mu \phi),$$

a continuous symmetry transformation of the fields leads, under suitable conditions, to a conserved current j^μ satisfying

$$\partial_\mu j^\mu = 0.$$

Integrating the time component over space gives a conserved charge,

$$Q = \int \mathrm{d}^3x \, j^0.$$

Definition 6.1: Introductory Noether statement

For a system whose action is invariant under a continuous global symmetry, there exists an associated conserved quantity. In field theory this appears as a conserved current and a conserved charge.

At MSc level it is worth seeing at least the skeleton of the derivation. Consider fields ϕ_i undergoing an infinitesimal transformation

$$\phi_i \to \phi_i + \delta\phi_i.$$

The first-order variation of the Lagrangian is

$$\delta\mathcal{L} = \sum_i \frac{\partial\mathcal{L}}{\partial\phi_i}\,\delta\phi_i + \sum_i \frac{\partial\mathcal{L}}{\partial(\partial_\mu\phi_i)}\,\delta(\partial_\mu\phi_i).$$

Using $\delta(\partial_\mu\phi_i) = \partial_\mu(\delta\phi_i)$ and integrating the second term by parts gives

$$\delta\mathcal{L} = \sum_i \left[\frac{\partial\mathcal{L}}{\partial\phi_i} - \partial_\mu\left(\frac{\partial\mathcal{L}}{\partial(\partial_\mu\phi_i)}\right)\right]\delta\phi_i + \partial_\mu\left(\sum_i \frac{\partial\mathcal{L}}{\partial(\partial_\mu\phi_i)}\,\delta\phi_i\right).$$

When the equations of motion hold, the bracketed term vanishes. If the transformation is a symmetry, then the change in the Lagrangian is at most a total derivative. In the simplest case of exact invariance, $\delta\mathcal{L} = 0$, and one obtains the conserved current

$$j^\mu = \sum_i \frac{\partial\mathcal{L}}{\partial(\partial_\mu\phi_i)}\,\delta\phi_i, \qquad \partial_\mu j^\mu = 0.$$

This is the core logic of Noether's theorem. The important thing for the student to notice is not only the final formula, but the chain of reasoning: continuous symmetry → controlled variation of the action → continuity equation → conserved charge.

6.2 Three standard examples

Three examples should become second nature early on:

- invariance under spatial translations leads to conservation of momentum,
- invariance under rotations leads to conservation of angular momentum,
- invariance under time translations leads to conservation of energy.

These are familiar from mechanics, but their meaning is broader. They show that symmetry governs what can change and what must remain fixed during time evolution.

Example 6.1: Global phase symmetry and charge conservation

Consider a complex scalar field with Lagrangian

$$\mathcal{L} = \partial_\mu \phi^* \, \partial^\mu \phi - m^2 \phi^* \phi.$$

Under a global phase transformation,

$$\phi \to e^{i\alpha}\phi, \qquad \phi^* \to e^{-i\alpha}\phi^*, \qquad \alpha = \text{constant},$$

the infinitesimal variations are

$$\delta\phi = i\alpha\phi, \qquad \delta\phi^* = -i\alpha\phi^*.$$

Substituting into the Noether formula gives

$$j^\mu = \frac{\partial \mathcal{L}}{\partial(\partial_\mu \phi)}\delta\phi + \frac{\partial \mathcal{L}}{\partial(\partial_\mu \phi^*)}\delta\phi^* = i\alpha\big(\phi^* \partial^\mu \phi - (\partial^\mu \phi^*)\phi\big).$$

Removing the common infinitesimal parameter, one obtains the conserved current

$$j^\mu = i\big(\phi^* \partial^\mu \phi - (\partial^\mu \phi^*)\phi\big), \qquad \partial_\mu j^\mu = 0.$$

In later gauge theory this is the seed of the connection between phase symmetry and electric charge.

6.3 Charges as generators

An especially important next step is that conserved charges often generate the symmetry transformation. In quantum mechanics and quantum field theory this is encoded by commutators. If Q is the conserved charge associated with a continuous symmetry, then the corresponding finite transformation may be written as

$$U(\alpha) = e^{i\alpha Q}.$$

Acting on an operator O gives

$$O \to O' = U(\alpha) O U(\alpha)^{-1}.$$

Expanding to first order in α yields

$$\delta O = O' - O = i\alpha \, [Q, O] + O(\alpha^2).$$

So the commutator with Q tells us how the operator changes under the symmetry. This is one of the places where the operator algebra of quantum mechanics and the symmetry language of field theory visibly become the same thing.

If Q generates a true symmetry of the dynamics, then it commutes with the Hamiltonian,

$$[H, Q] = 0.$$

Using the Heisenberg equation of motion,

$$\frac{dQ}{dt} = \frac{i}{\hbar} \, [H, Q] + \left(\frac{\partial Q}{\partial t}\right)_{\text{explicit}},$$

one sees that a time-independent symmetry generator is conserved. This is the operator-language counterpart of the continuity equation derived above.

Take-home message

Noether's theorem teaches a central lesson for the Standard Model: symmetry is not merely a way of naming states. It shapes dynamics, produces conserved charges, and prepares the ground for gauge structure.

7 Groups: the language of transformations

7.1 Why transformations need mathematical organisation

If symmetries are to play a structural role in building a theory, the set of all symmetry transformations must be described precisely. That description is group theory.

Definition 7.1: Group

A *group* G is a set equipped with a composition law such that:

(i) closure: if $g, h \in G$, then $gh \in G$,

(ii) identity: there exists $e \in G$ such that $eg = ge = g$ for all $g \in G$,

(iii) inverse: for each $g \in G$ there exists g^{-1} with $gg^{-1} = g^{-1}g = e$,

(iv) associativity: $(gh)k = g(hk)$ for all $g, h, k \in G$.

If in addition $gh = hg$ for all $g, h \in G$, the group is called *Abelian*.

Each axiom has physical meaning. Closure says that performing two allowed symmetry transformations in succession yields another allowed transformation. The identity means that doing nothing is allowed. Inverses mean one can undo a symmetry operation. Associativity guarantees consistency of repeated composition.

7.2 Elementary examples

Example 7.1: Phase rotations and $\mathrm{U}(1)$

The transformations

$$\psi \to \mathrm{e}^{\mathrm{i}\alpha}\psi, \qquad \alpha \in \mathbb{R},$$

form the group $\mathrm{U}(1)$ under multiplication. It is Abelian because

$$\mathrm{e}^{\mathrm{i}\alpha}\mathrm{e}^{\mathrm{i}\beta} = \mathrm{e}^{\mathrm{i}\beta}\mathrm{e}^{\mathrm{i}\alpha}.$$

Example 7.2: Rotations in the plane and $\mathrm{SO}(2)$

Rotations in two-dimensional Euclidean space form the group $\mathrm{SO}(2)$. The group operation is composition of rotations, and the inverse of a rotation by angle θ is a rotation by $-\theta$.

Example 7.3: Rotations in three dimensions and SO(3)

Rotations in three-dimensional space form the group SO(3). This group is non-Abelian: a rotation about the x axis followed by one about the y axis is not, in general, the same as performing them in the opposite order.

7.3 Why U(1), SU(2), and SU(3) matter

The Standard Model is built around the internal gauge structure

$$\mathrm{SU}(3)_c \times \mathrm{SU}(2)_L \times \mathrm{U}(1)_Y.$$

This already tells us that groups are not peripheral mathematical curiosities. They are the organising framework for the interactions of the theory. Tong's notes emphasise this strongly: much of the Standard Model follows from gauge invariance together with the representation content of matter under these groups. Module 1 therefore prepares the student to see group theory not as an abstract appendix, but as the mathematical home of the Standard Model interactions.

7.4 A small dictionary of common matrix groups

Before moving on, it is helpful to keep a compact dictionary of the groups that most often appear in a first course on particle physics.

Group	Informal description	Why it matters in physics
U(1)	Complex phases of unit modulus	Simplest Abelian symmetry; prototype for electromagnetism and hypercharge language
SO(2)	Rotations in a plane	Simplest continuous geometric rotation group
SO(3)	Rotations in three-dimensional space	Classical rotations and angular momentum structure
SU(2)	2×2 unitary matrices with determinant one	Spin one-half, weak isospin, simplest non-Abelian example
SU(3)	3×3 unitary matrices with determinant one	Colour symmetry in QCD and the octet structure of gauge bosons

The purpose of this table is not technical completeness. It is to help the student recognise that the familiar groups of quantum mechanics and the groups of the Standard Model belong to the same mathematical family.

8 Continuous groups, generators, and Lie algebras

8.1 Infinitesimal transformations and exponentiation

For continuous symmetries it is enough to understand infinitesimal transformations; finite transformations are then built by exponentiation. This is the central idea behind Lie groups and Lie algebras.

Definition 8.1: Generator of a continuous transformation

If a continuous symmetry depends on parameters α^a, a finite transformation may be written schematically as

$$U(\alpha) = \exp(\mathrm{i}\alpha^a T^a)\,,$$

where the T^a are the generators.

For a one-parameter symmetry this becomes simply

$$U(\alpha) = \exp(\mathrm{i}\alpha T).$$

The parameter tells us how far we move along the symmetry, while the generator tells us in which direction in symmetry space we move.

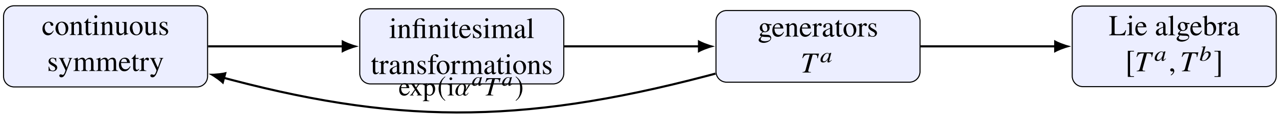

Figure 5: Conceptual flow from continuous symmetry to its Lie-algebraic description. A continuous symmetry is understood locally through its generators and globally through exponentiation.

8.2 A more explicit look at generators

For students seeing this material for the first time, the generator can feel mysterious. The key idea is simply that the generator captures the infinitesimal version of the transformation. Consider a one-parameter family of transformations $U(\alpha)$. If the family is smooth, then near the identity one may write

$$U(\alpha) = 1 + \mathrm{i}\alpha T + O(\alpha^2).$$

The object T is the generator. Knowing it is enough to reconstruct the finite transformation through exponentiation.

This is analogous to ordinary calculus. A derivative at one point contains local information about how a function changes. In Lie theory, the generator contains local information about how the group transformation changes near the identity. This is why so much of the analysis of continuous symmetry can be reduced to the study of generators and their commutators.

One can make this idea concrete by differentiating at the identity:

$$T = -\mathrm{i}\,\frac{\mathrm{d}U(\alpha)}{\mathrm{d}\alpha}\bigg|_{\alpha=0}.$$

So the generator is literally the first derivative of the group element with respect to its continuous parameter, up to a conventional factor of $-\mathrm{i}$. In this sense the Lie algebra is the tangent-space data of the Lie group near the identity.

Example 8.1: Rotations as accumulated infinitesimal steps

A finite spatial rotation may be understood as the accumulation of many infinitesimal rotations. The exponential form

$$U(\theta) = \exp\left(-\frac{\mathrm{i}}{\hbar}\theta J\right)$$

encodes exactly that idea. Here J is the angular-momentum generator of the rotation. Expanding for small θ gives

$$U(\theta) = 1 - \frac{\mathrm{i}}{\hbar}\theta J + O(\theta^2),$$

so J tells us how the state begins to change under an infinitesimal rotation.

8.3 Lie algebras and commutation relations

The generators satisfy commutation relations of the form

$$\left[T^a, T^b\right] = \mathrm{i} f^{abc} T^c,$$

where the constants f^{abc} are called structure constants. These commutation relations define the corresponding Lie algebra.

A useful way to understand why commutators appear is to compose two infinitesimal transformations. If

$$U(\alpha) = 1 + \mathrm{i}\alpha^a T^a, \qquad U(\beta) = 1 + \mathrm{i}\beta^b T^b,$$

then the difference between applying them in opposite orders is

$$U(\alpha)U(\beta) - U(\beta)U(\alpha) = -\alpha^a \beta^b \left[T^a, T^b\right] + O(\alpha^2\beta, \alpha\beta^2).$$

For the set of transformations to close, this difference must itself correspond to another infinitesimal transformation. That is only possible if the commutator of two generators is again a linear combination of generators. This is the operational meaning of the Lie algebra.

Definition 8.2: Lie algebra

A *Lie algebra* is a vector space equipped with a bilinear antisymmetric bracket satisfying the Jacobi identity. In the matrix realisations relevant to physics, the bracket is the commutator.

The Jacobi identity reads

$$[T^a, [T^b, T^c]] + [T^b, [T^c, T^a]] + [T^c, [T^a, T^b]] = 0.$$

This ensures internal consistency of the algebraic structure.

8.4 Abelian versus non-Abelian algebras

For an Abelian group such as U(1), all generators commute, so the structure constants vanish. For non-Abelian groups such as SU(2) and SU(3), the generators do not commute. This distinction is not just mathematical. In gauge theory it produces very different dynamics. Abelian gauge fields do not self-interact in the same basic way as non-Abelian gauge fields do.

Example 8.2: The algebra of SU(2)

A convenient basis for the generators of SU(2) is

$$T^a = \frac{\sigma^a}{2},$$

where σ^a are the Pauli matrices. They satisfy

$$[T^a, T^b] = \mathrm{i}\epsilon^{abc}T^c.$$

This is the simplest non-Abelian Lie algebra students encounter repeatedly in quantum mechanics and particle physics.

Example 8.3: Rotations and angular momentum

The generators of spatial rotations are the angular momentum operators J_i, with algebra

$$[J_i, J_j] = \mathrm{i}\hbar\epsilon_{ijk}J_k.$$

This relation is both an operator statement in quantum mechanics and the Lie algebra of spatial rotations. The same mathematical object is doing both jobs.

8.5 Why this matters later

The Standard Model gauge groups are Lie groups, and their local structure is encoded in Lie algebras. If students later meet Yang–Mills fields, covariant derivatives, or field strengths without first understanding what a generator is or why commutators appear, the theory will look arbitrary. Module 1 aims to make those later structures feel inevitable rather than surprising.

9 Representations and multiplets

9.1 Why groups become physical only through representations

A group by itself is an abstract set of transformations. To become physically meaningful, we must specify what those transformations act on and how they act. This is the role of representation theory.

Definition 9.1: Representation

A *representation* of a group G on a vector space V is a map assigning to each element $g \in G$ a linear transformation $D(g)$ on V such that

$$D(g_1g_2) = D(g_1)D(g_2).$$

The same symmetry group can therefore act differently on different physical objects. Some may be unchanged, some may mix in pairs, some in triples, and so on. The group is the same, but the representation differs.

9.2 Scalars, vectors, spinors, and internal multiplets

At this stage it is useful to distinguish two uses of the word representation. Physical objects transform under spacetime symmetries, leading to scalars, vectors, and spinors. They also transform under internal symmetries, leading to singlets, doublets, triplets, and higher multiplets. A field may simultaneously carry both kinds of labels: for example, a Lorentz spinor, a weak doublet, and a color singlet.

The key lesson is that a representation tells us how components mix under the symmetry. A singlet remains unchanged. A doublet mixes two components. A triplet mixes three. For an infinitesimal SU(2) transformation a doublet

$$\Psi = \begin{pmatrix} \psi_1 \\ \psi_2 \end{pmatrix}$$

transforms as

$$\Psi \to \Psi' = \left(\mathbf{1} + i\alpha^a \frac{\sigma^a}{2} \right) \Psi + O(\alpha^2).$$

So the symmetry does not merely relabel the field; it acts linearly on the vector of components. That is what it means to say that the field belongs to a doublet representation.

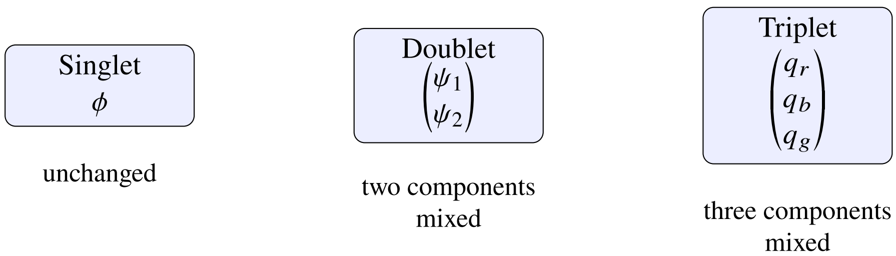

Figure 6: Schematic representation sketch. The same symmetry can act trivially on a singlet or mix the components of a doublet or triplet.

9.3 First Standard Model examples

Already at introductory level one can preview how representation theory organises the Standard Model:

- quarks transform as color triplets under $SU(3)_c$,
- left-handed fermions appear in weak doublets under $SU(2)_L$,
- many fields are singlets under one or more factors,
- hypercharge labels how fields transform under $U(1)_Y$.

Example 9.1: A first Standard Model multiplet preview

The left-handed electron and electron neutrino are grouped into a weak doublet,

$$L_e = \begin{pmatrix} \nu_e \\ e \end{pmatrix}_L .$$

This notation already tells us something structural: the two fields transform together under $SU(2)_L$. By contrast, the right-handed electron e_R is a weak singlet. Hence the weak interaction treats left- and right-handed fields differently.

At the level of Module 1, the purpose of such examples is not to study chirality in detail, but to show that representation theory is the mechanism that turns symmetry groups into physical particle assignments.

9.4 Representation labels as physical information

When students first encounter labels such as singlet, doublet, or triplet, they can sound like mere classification tags. But in particle physics they carry direct physical meaning. They tell us how many components of a field are mixed by the symmetry and therefore which interaction structures are possible. A weak doublet couples to the weak gauge bosons in a way that a weak singlet does not. A color triplet interacts with gluons in a way that a color singlet does not. Representation theory is therefore not an optional refinement. It is the mechanism that turns group theory into particle physics.

It is also worth noting that representation labels are layered. A field can be a Lorentz scalar yet a weak doublet. Another field can be a Lorentz spinor yet a color singlet. The full identity of a Standard Model field is really a list of representation labels under several symmetry groups at once.

10 Invariants and the logic of Lagrangian building

10.1 Why a Lagrangian is not arbitrary

In modern field theory dynamics is encoded in a Lagrangian density $\mathcal{L}$. But the Lagrangian is not a random expression. It must be assembled from terms whose combined transformation law respects the symmetries of the theory. If the action

$$S = \int d^4x\, \mathcal{L}$$

is required to be invariant, then only invariant or appropriately covariant combinations of fields and derivatives may appear.

Definition 10.1: Invariant term in the Lagrangian

A term in the Lagrangian is called *invariant* under a symmetry if its total transformation leaves the action unchanged. In practice this strongly restricts the allowed combinations of fields.

This is one of the most powerful ideas in particle physics. Once the symmetry group and the field content are specified, many candidate terms are immediately ruled out.

10.2 A schematic example

Consider a complex scalar field with global phase symmetry

$$\phi \to e^{i\alpha}\phi.$$

Then its complex conjugate transforms as

$$\phi^* \to e^{-i\alpha}\phi^*.$$

The combination $\phi^*\phi$ is therefore invariant:

$$\phi^*\phi \to e^{-i\alpha}\phi^* e^{i\alpha}\phi = \phi^*\phi.$$

A mass term of the form $m^2\phi^*\phi$ is allowed by the symmetry, whereas a linear term proportional to ϕ would break it.

The same logic applies to derivatives. Because the parameter is global here, one has

$$\partial_\mu\phi \to e^{i\alpha}\partial_\mu\phi, \qquad \partial_\mu\phi^* \to e^{-i\alpha}\partial_\mu\phi^*,$$

so the kinetic term is also invariant:

$$\partial_\mu\phi^*\,\partial^\mu\phi \to e^{-i\alpha}\partial_\mu\phi^*\,e^{i\alpha}\partial^\mu\phi = \partial_\mu\phi^*\,\partial^\mu\phi.$$

Thus the symmetry simultaneously constrains mass terms and kinetic terms. The broader lesson is that invariance is checked term by term, not only at the level of vague physical intuition.

A similar idea later appears for fermions. A bilinear such as $\bar{\psi}\psi$ can be invariant under some symmetries and not under others, depending on how the field transforms. This is why the representation content of the fields matters so much when one builds a Lagrangian.

Example 10.1: Symmetry as an exclusion principle

A useful way to think about symmetry is as an exclusion principle for theory-building. Symmetry does not merely add elegance. It forbids incorrect terms. This is one reason the Standard Model is powerful: once the symmetry group and field content are fixed, the theory is much more rigid than a generic Lagrangian would be.

10.3 Why this matters for the Standard Model

The Standard Model is highly constrained because its fields are placed in representations of $SU(3)_c \times SU(2)_L \times U(1)_Y$ and the action must respect the corresponding local symmetry. Long before we write the full Standard Model Lagrangian, Module 1 should already make students comfortable with the idea that symmetry excludes many seemingly possible terms and therefore acts as a constructive principle.

11 From global symmetry to local symmetry

11.1 Global symmetry

A global symmetry acts with the same transformation everywhere in spacetime. Schematically,

$$\psi(x) \to U\psi(x),$$

with U independent of x. Such symmetries are the natural domain of the introductory Noether theorem.

11.2 Local symmetry and the derivative problem

A local symmetry allows the transformation to depend on spacetime position:

$$\psi(x) \longrightarrow U(x)\psi(x).$$

At first sight this looks like a small change. In fact it changes everything. The ordinary derivative no longer transforms covariantly:

$$\partial_\mu\psi(x) \longrightarrow \partial_\mu\left(U(x)\psi(x)\right) = U(x)\partial_\mu\psi(x) + \left(\partial_\mu U(x)\right)\psi(x).$$

The extra term involving $\partial_\mu U(x)$ destroys the naive symmetry. This is the derivative problem: once the transformation parameter varies from point to point, neighboring spacetime points can no longer be compared by the ordinary derivative in a symmetry-respecting way.

For the simplest Abelian case, one may write

$$\psi(x) \longrightarrow \psi'(x) = \mathrm{e}^{\mathrm{i}q\alpha(x)}\psi(x).$$

Then

$$\partial_\mu\psi'(x) = \mathrm{e}^{\mathrm{i}q\alpha(x)}\left(\partial_\mu\psi + \mathrm{i}q\,\partial_\mu\alpha\,\psi\right).$$

The troublesome term is the one proportional to $\partial_\mu\alpha$. It vanishes for global symmetry because α is constant, but it survives for local symmetry. That is why the local case is qualitatively different from the global one.

Definition 11.1: Covariant derivative: schematic form

To restore local symmetry one introduces a gauge field and defines the covariant derivative. In a convention consistent with Modules 2 and 3,

$$D_\mu = \partial_\mu - \mathrm{i}qA_\mu$$

in the Abelian case, and more generally,

$$D_\mu = \partial_\mu - \mathrm{i}gA_\mu^a T^a$$

for a non-Abelian group. The transformation law of the gauge field is chosen so that $D_\mu\psi$ transforms covariantly when ψ does.

In the Abelian example above, the required gauge-field transformation is

$$A_\mu \longrightarrow A'_\mu = A_\mu + \partial_\mu\alpha,$$

if the matter field transforms as $\psi \longrightarrow \mathrm{e}^{\mathrm{i}q\alpha(x)}\psi$ and one uses $D_\mu = \partial_\mu - \mathrm{i}qA_\mu$. A short substitution then shows

$$D'_\mu\psi' = \mathrm{e}^{\mathrm{i}q\alpha(x)}D_\mu\psi,$$

so the covariant derivative transforms in the same way as the field itself. This is the precise sense in which the gauge field repairs the failure of the ordinary derivative.

This is the conceptual birth of interactions in gauge theory. Forces are not inserted into the theory by hand. They arise because local symmetry requires the introduction of gauge fields.

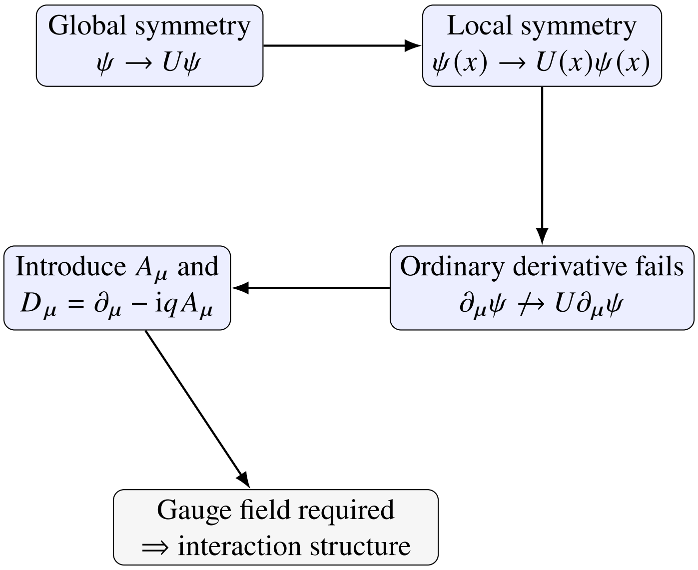

Figure 7: Global versus local symmetry. The failure of the ordinary derivative under local transformations forces the introduction of a gauge field and a covariant derivative.

11.3 Abelian and non-Abelian gauge structure

For an Abelian symmetry like U(1) this logic leads to electrodynamics. For non-Abelian symmetries like SU(2) and SU(3) it leads to Yang–Mills theories, where the gauge fields carry group indices and the Lie algebra enters directly into the dynamics. At this stage the key lesson is conceptual: local symmetry is far more constraining than global symmetry, and its implementation is what leads naturally to gauge interactions.

Take-home message

The statement *local symmetry requires gauge fields* is one of the central conceptual outputs of Module 1. It is the bridge from formal language to the Standard Model itself.

11.4 Why gauge fields are not optional additions

A common beginner's misunderstanding is to imagine that one first writes a matter theory and then later if desired, add a force field by hand. Gauge theory turns this picture upside down. Once one insists on local symmetry, the gauge field is not optional. It is required by consistency. In this sense gauge interactions are structural rather than decorative.

This is one of the reasons the gauge principle has such explanatory power. It does not merely tell us that interactions exist. It tells us why they enter the theory in a highly specific way. The covariant derivative is not guessed. It is forced.

11.5 Abelian versus non-Abelian field content

In an Abelian theory the gauge parameter is effectively a single function of spacetime, and the associated gauge field is correspondingly simple. In a non-Abelian theory there is one gauge field component for each generator of the Lie algebra. Thus a theory based on SU(2) carries three gauge-field components, while one based on SU(3) carries eight. This is the first conceptual reason why the weak interaction has three gauge bosons before symmetry breaking and why the strong interaction has eight gluons.

At this stage one need not derive the field strength or the Yang–Mills Lagrangian in detail. It is enough to see that the number and structure of gauge fields are already encoded by the symmetry group itself.

12 First roadmap to the Standard Model

We are now in a position to see how the ideas of Module 1 point toward the Standard Model without yet constructing the full theory.

12.1 The toolkit and its destination

Tool from Module 1	What it later controls in the Standard Model
Hilbert space and operators	Quantum states, field operators, currents, charges, observables
Commutators	Compatibility, uncertainty, symmetry-generator algebra
Symmetry and invariance	Conservation laws, selection rules, admissible dynamics
Groups and Lie algebras	Internal gauge groups $\mathrm{U}(1)_Y$, $\mathrm{SU}(2)_L$, $\mathrm{SU}(3)_c$
Representations and multiplets	Assignments of fermions and bosons to singlets, doublets, triplets, adjoints
Invariant combinations	Allowed terms in the Lagrangian
Local symmetry	Gauge fields and the interaction principle itself

12.2 Why the gauge group looks like a product

At introductory level we do not yet derive why nature chooses exactly $\mathrm{U}(1)_Y$, $\mathrm{SU}(2)_L$, and $\mathrm{SU}(3)_c$. But we can already see why the appearance of a direct product of Lie groups is natural. Different interaction sectors are associated with different internal symmetries. Matter fields are assigned to representations of the product group, and those assignments determine their transformation properties and allowed couplings.

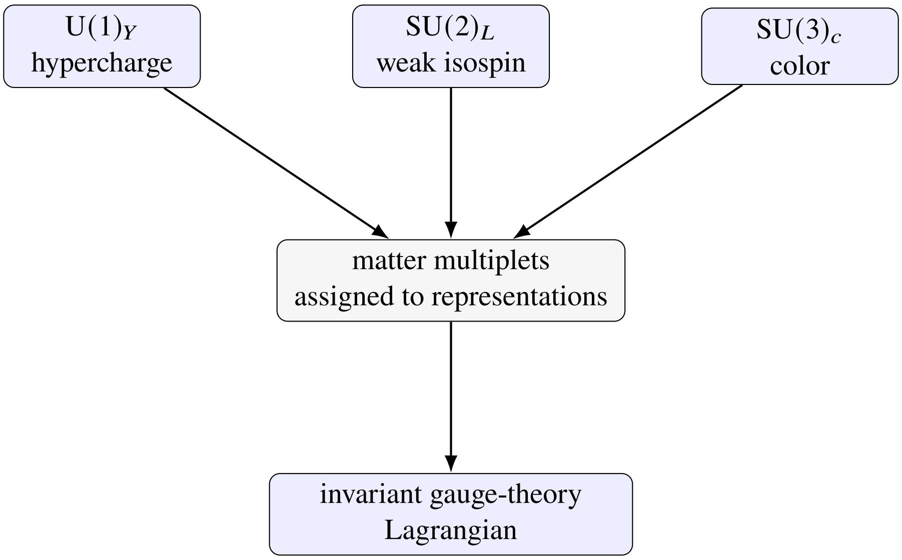

Figure 8: A first roadmap to the Standard Model. Group factors, representations, and invariance together organise the structure of the theory.

13 Bridge to Module 2 and Module 3

13.1 Why Module 2 comes next

Module 1 has deliberately treated states, operators, symmetry, groups, and invariants in a way that is not yet fully relativistic. But the Standard Model is a relativistic quantum field theory. The next natural question is therefore: how must physical objects transform under spacetime symmetry? This leads directly to Module 2, where students study Lorentz transformations, four-vectors, spinors, Dirac matrices, and the Dirac equation.

The course structure states this explicitly. Module 2 is the relativistic framework for elementary particles and is meant to explain why spinors are required and how relativistic invariance constrains the construction of particle-physics models. Module 1 has therefore prepared the abstract language, while Module 2 will supply the specific spacetime representations needed for relativistic matter fields.

13.2 Why Module 3 follows naturally

At the end of Module 1 we have already encountered the distinction between global and local symmetry and seen that local symmetry forces the introduction of gauge fields. This is the seed of Module 3. There the general idea becomes concrete: first through quantum electrodynamics as the Abelian prototype, then through non-Abelian gauge theory, and finally through the full gauge structure

$$\mathrm{SU}(3)_c \times \mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$$

of the Standard Model.

The course structure also emphasises that Module 3 should show the logical construction of the Standard Model from symmetry requirements, field content, and local gauge invariance. This is exactly why Module 1 spends time on groups, generators, representations, and local symmetry. Without them, the construction of Module 3 would appear from nowhere.

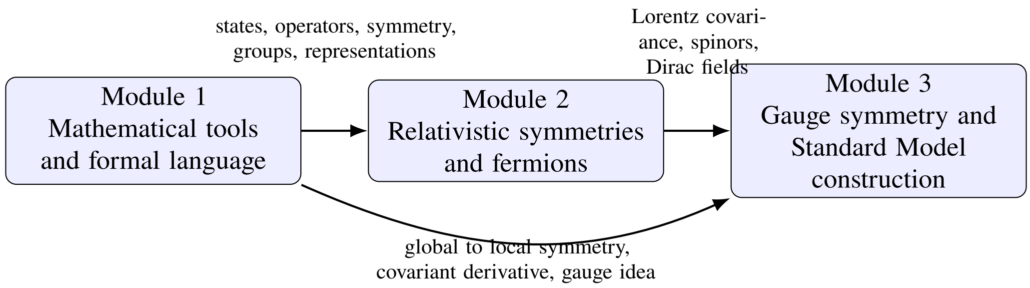

Figure 9: Final conceptual map connecting Module 1 to Module 2 and Module 3. Module 1 provides the formal grammar, Module 2 develops the relativistic spacetime framework, and Module 3 uses this full language to construct gauge theory and the Standard Model.

14 Common conceptual confusions and how to avoid them

Because Module 1 introduces a large amount of formal language quickly, it is useful to end with a short list of common confusions.

14.1 A state is not an observable

A state tells us how the system is prepared. An observable tells us what may be measured. The two are not interchangeable. Mixing them up causes confusion later when one moves to operator-valued fields.

14.2 A symmetry is not the same as an invariant

The symmetry is the transformation. The invariant is the quantity left unchanged by that transformation. Keeping this distinction clear helps when moving from intuitive examples to Lagrangian building.

14.3 A group is not yet a representation

The group is the abstract structure of transformations. A representation tells us how those transformations act on specific physical objects. Without a representation, the symmetry has not yet been connected to a field or particle multiplet.

14.4 Global and local symmetry are conceptually different

A local symmetry is not obtained by casually replacing a constant parameter with a function of spacetime and hoping the theory remains unchanged. The derivative term changes, and the need to repair that failure is what introduces the gauge field. This difference is so important that students should revisit it several times.

14.5 Why this list matters

These confusions are not trivial. Each one corresponds to a real conceptual obstacle students meet when they first encounter relativistic field theory and gauge structure. Writing them down explicitly is therefore part of the supporting role of these notes.

15 Summary and take-home messages

These notes have expanded the logic of the Module 1 slides into a self-contained text. The central question of the module has been: what is the minimal mathematical language needed to construct the Standard Model in a logically coherent way? The answer may now be summarised as follows.

1. Quantum theory begins with states, operators, eigenvalues, and expectation values.
2. Operator algebra matters because commutators encode compatibility, uncertainty, and eventually symmetry-generator structure.
3. A symmetry is a transformation that leaves the physical content unchanged; an invariant is a quantity preserved by that transformation.
4. Continuous global symmetries lead to conserved currents and charges through Noether's theorem.
5. Groups and Lie algebras provide the mathematical language of continuous transformations.
6. Representations tell us how fields and particles transform, and therefore how they are organised into multiplets.
7. The Lagrangian is not arbitrary; it is built from invariant combinations of fields.
8. The jump from global to local symmetry forces gauge fields and creates the conceptual origin of interactions.

At the close of Module 1, students should not yet feel that they have fully constructed the Standard Model. But they should feel something more important: that the later construction will no longer look mysterious. Once one sees why operator language, symmetry, group theory, and local invariance are indispensable, the Standard Model begins to appear not as a disconnected pile of formulas, but as a mathematically organised framework whose structure can be understood step by step.

Take-home message

The Standard Model is built from a remarkably small toolkit: state space, operators, commutators, symmetry, invariance, groups, generators, Lie algebras, representations, and the local gauge principle. Module 1 provides exactly this toolkit.

Recommended reading

The following sources are particularly helpful for strengthening the material of this module before moving to Modules 2 and 3.

- Mark Thomson, *Modern Particle Physics*, especially Chapter 9 on symmetries and the quark model, Section 10.1 on the local gauge principle, and Appendix E on Noether's theorem.
- David Tong, *The Standard Model*, especially the early material on symmetries, gauge invariance, and Lie algebras.
- Fernando Quevedo, *The Standard Model*, especially the sections on internal symmetries, Noether's theorem, charges as generators, and the origin of gauge symmetry.

Compact notation sheet for Module 1

Symbol	**Meaning**
$\lvert\psi\rangle$	state vector
$\langle\psi\rvert$	dual vector
$\hat{A}$	linear operator / observable
$\hat{A}\lvert a\rangle = a\lvert a\rangle$	eigenvalue equation
$\langle\hat{A}\rangle_\psi$	expectation value of $\hat{A}$ in the state $\lvert\psi\rangle$
$\mathrm{Var}(A) = \langle\hat{A}^2\rangle_\psi - \langle\hat{A}\rangle_\psi^2$	variance of the observable A
$\left[\hat{A}, \hat{B}\right]$	commutator of two operators
$U(\alpha) = \exp(\mathrm{i}\alpha^a T^a)$	finite continuous transformation
T^a	generators of a continuous symmetry
$\left[T^a, T^b\right] = \mathrm{i}f^{abc}T^c$	Lie algebra relation
$D_\mu = \partial_\mu - \mathrm{i}qA_\mu$	Abelian covariant derivative
$D_\mu = \partial_\mu - \mathrm{i}gA_\mu^a T^a$	non-Abelian covariant derivative
$\mathrm{SU}(3)_c \times \mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$	Standard Model gauge structure

About these notes

These notes are intended as supporting lecture notes for Module 1 on the mathematical tools and formal language needed for a first systematic study of the Standard Model. Their aim is to provide a clear and self-contained introduction to the concepts that recur throughout modern particle physics, including states, operators, commutators, symmetries, groups, generators, representations, and invariants. The emphasis is on conceptual clarity, mathematical precision, and the connection between formal structure and physical meaning. In particular, the notes are designed to show how the language introduced here prepares the ground for later topics such as relativistic fields, Lorentz symmetry, gauge invariance, and the construction of the Standard Model. These notes are not intended to replace more detailed textbooks. Rather, they are meant to serve as a bridge between undergraduate quantum mechanics and the more advanced study of quantum field theory and particle physics.