1 Introduction

Scattering experiments are crucial for our understanding of the building blocks of nature. The standard model of particle physics was developed from scattering experiments, including the discovery of the weak force bosons $W^{\pm}$ and Z^0, quarks and gluons, and most recently the Higgs boson.

The key observable measured in particle scattering experiments is the ***scattering cross-section*** σ. It encodes the likelihood of a given process to take place as a function of the energy and momentum of the particles involved. A more refined version of this quantity is the ***differential cross-section*** $d\sigma/d\Omega$: it describes the dependence of the cross-section on the angles of the scattered particles.

Interpretation of data from scattering experiments relies heavily on theoretical predictions of scattering cross-sections. These are calculated in ***relativistic quantum field theory (QFT)***, which is the mathematical language for describing elementary particles and their interactions. Relativistic QFT combines special relativity with quantum physics and is a hugely successful and experimentally well-tested framework for describing elementary particles and the fundamental forces of nature. In quantum mechanics, the probability distribution $|\psi|^2 = \psi^*\psi$ for a particle is given by the norm-squared of its complex-valued wavefunction ψ. Analogously, in quantum field theory, the differential cross-section is proportional to the norm-squared of the ***scattering amplitude*** A, $d\sigma/d\Omega \propto |A|^2$. The amplitudes A are well-defined physical observables: they are the subject of this book.

Scattering amplitudes have physical relevance through their role in the scattering cross-section. Moreover, it has been realized in recent years that amplitudes themselves have a very interesting mathematical structure. Understanding this structure guides us towards more efficient methods to calculate amplitudes. It also makes it exciting to study scattering amplitudes in their own right and explore (and exploit) their connections to interesting branches of mathematics, including combinatorics and geometry.

The purpose of this book is two-fold. First, we wish to provide a pedagogical introduction to the efficient modern methods for calculating scattering amplitudes. Second, we survey several interesting mathematical properties of amplitudes and recent research advances. The hope is to bridge the gap between a standard graduate course in quantum field theory and the modern approach to scattering amplitudes as well as current research. Thus we assume readers to be familiar with the basics of quantum field theory and particle physics; however, here in the first chapter we offer an introduction that may also be helpful for a broader audience with an interest in the subject.

What is a scattering amplitude?

Let us consider a few examples of scattering processes involving electrons e^-, positrons e^+, and photons γ:

$$\begin{aligned} \text{Compton scattering} \quad & e^- + \gamma \ \to e^- + \gamma\,, \\ \text{Møller scattering} \quad & e^- + e^- \to e^- + e^-\,, \\ \text{Bhabha scattering} \quad & e^- + e^+ \to e^- + e^+\,, \\ e^-e^+ \text{ annihilation} \quad & e^- + e^+ \to \gamma + \gamma\,. \end{aligned} \tag{1.1}$$

These processes are described in the quantum field theory that couples Maxwell's electromagnetism to electrons and positrons, namely quantum electrodynamics (QED). Each process is characterized by the types of particles involved in the initial and final states as well as the relativistic momentum $\vec{p}$ and energy E of each particle. This input is the ***external data*** for an amplitude: a scattering amplitude A_n involving a total of n initial and final state particles takes the list $\{E_i, \vec{p}_i\,;\, \text{type}_i\}$, $i = 1, 2, \ldots, n$, of external data and returns a complex number:

$$n\text{-particle amplitude } A_n: \qquad \big\{E_i, \vec{p}_i\,;\, \text{type}_i\big\} \to A_n\Big(\big\{E_i, \vec{p}_i\,;\, \text{type}_i\big\}\Big) \in \mathbb{C}\,. \tag{1.2}$$

"Particle type" involves more than saying which particles scatter: it also includes a specification of the appropriate quantum numbers of the initial and final states, for example the polarization of a photon or the spin-state of a fermion.

In special relativity, the energy E and momentum $\vec{p}$ of a physical particle must satisfy the relationship $E^2 = |\vec{p}|^2c^2 + m^2c^4$ with m the rest mass of the particle. This is simply the statement that $E^2 - |\vec{p}|^2c^2$ is a Lorentz-invariant quantity in Minkowski space. Thus, one of the constraints on the external data for an n-particle amplitude is that the relativistic ***on-shell*** condition,

$$E_i^2 - |\vec{p}_i|^2c^2 = m_i^2c^4\,, \tag{1.3}$$

holds for all particles $i = 1, \ldots, n$. The initial and final state electrons and positrons in the QED processes (1.1) must satisfy the on-shell condition (1.3) with m_i equal to the electron mass, $m_e = 511\,\text{keV}/c^2$, while m_i is zero for the photon since it is massless.

It is convenient to work in "natural units" where the speed of light c and Planck's constant $\hbar = h/2\pi$ are set to unity: $c = 1 = \hbar$. Proper units can always be restored by dimensional analysis. We combine the energy and momentum into a 4-momentum $p_i^\mu = (E_i, \vec{p}_i)$, with $\mu = 0, 1, 2, 3$, and write $p_i^2 = -E_i^2 + |\vec{p}_i|^2$ such that the on-shell condition (1.3) becomes

$$p_i^2 = -m_i^2\,. \tag{1.4}$$

Conservation of relativistic energy and momentum requires the sum of initial momenta p^μ_{in} to equal the sum of final state momenta p^μ_{out}. It is convenient to flip the signs of all incoming momenta so that the ***conservation of 4-momentum*** simply reads

$$\sum_{i=1}^{n} p_i^\mu = 0 . \tag{1.5}$$

Thus, to summarize, the external data for the amplitude involve a specification of 4-momentum and particle type for each external particle, $\{p_i^\mu, \text{type}_i\}$, subject to the on-shell constraints $p_i^2 = -m_i^2$ and momentum conservation (1.5). We often use the phrase ***on-shell amplitude*** to emphasize that the external data satisfy the kinematic constraints (1.4) and (1.5) and include the appropriate polarization vectors or fermion spin wavefunctions.

Feynman diagrams

Scattering amplitudes are typically calculated as a perturbation series in the expansion of a small dimensionless parameter that encodes the coupling of interactions between the particles. For QED processes, this dimensionless coupling is the fine structure constant[1] $\alpha \approx 1/137$. In the late 1940s, Feynman introduced a diagrammatic way to organize the perturbative calculation of the scattering amplitudes. It expresses the n-particle scattering amplitude A_n as a sum of all possible Feynman diagrams with n external legs:

$$A_n = \sum (\text{Feynman diagrams}) =$$

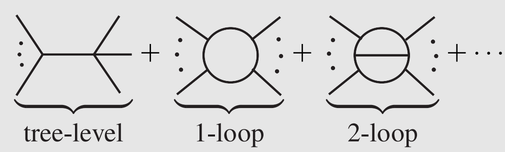

(1.6)

The sum of diagrams is organized by the number of closed loops. For a given number of particles n, the loop-diagrams are suppressed by powers of the coupling, e.g. in QED an L-loop diagram is order α^{2L} compared with the tree-level. Hence the loop-expansion is a diagrammatic representation of the perturbation expansion. The leading contribution comes from the tree-diagrams; their sum is called the ***tree-level amplitude***, A_n^{tree}. The next order is 1-loop; the sum of the 1-loop diagrams is the ***1-loop amplitude*** $A_n^{\text{1-loop}}$, etc. The full amplitude is then

$$A_n = A_n^{\text{tree}} + A_n^{\text{1-loop}} + A_n^{\text{2-loop}} + \cdots \tag{1.7}$$

It is rare that the loop-expansion is convergent; typically, the number of diagrams grows so quickly at higher-loop order that it overcomes the suppression from the higher powers in the small coupling. Amplitudes can be Borel summable, but this is not a subject we treat here. Instead, we pursue an understanding of the contributions to the amplitude order-by-order in perturbation theory.

[1] In SI units, the fine structure constant is $\alpha = e^2/(4\pi\epsilon_0\hbar c)$, where ϵ_0 is the vacuum permittivity.

The quantum field theories we are concerned with in this book are defined by Lagrangians that encode the particles and how they interact. The Lagrangian for QED determines a simple basic 3-particle interaction between electrons/positrons and the photon:

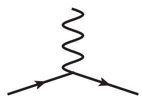

. (1.8)

From this, one can build diagrams such as

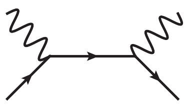

, 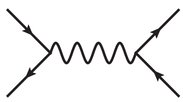 , and 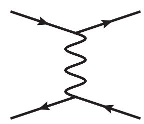 . (1.9)

Reading the diagrams left to right, the first diagram describes the absorption and subsequent emission of a photon (wavy line) by an electron (solid line), and as such it contributes at the leading order in perturbation theory to the Compton scattering process $e^- + \gamma \;\rightarrow\; e^- + \gamma$. The second and third diagrams in (1.9) are the two tree-level diagrams that give the leading-order contribution to the Bhabha scattering process $e^- + e^+ \;\rightarrow\; e^- + e^+$. The first of those two diagrams encodes the annihilation of an electron and a positron to a photon and the subsequent e^+e^- pair-creation. The second diagram is the exchange of a photon in the scattering of an electron and a positron.

A Feynman diagram is translated to a mathematical expression via Feynman rules. These rules are specific to the particle types and the theory that describes their interactions. A Feynman graph has the following essential parts:

- *External lines:* the Feynman diagram has an external line for each initial and final state particle. The rule for the external line depends on the particle type. For spin-1 particles, such as photons and gluons, the external line rule encodes the polarizations. For fermions, it contains information about the spin state.
- *Momentum labels:* every line in the Feynman diagram is associated with a momentum. For the external lines, the momentum is fixed by the external data. Momentum conservation is enforced at any vertex. For a tree-graph, this fully determines the momenta on all internal lines. For an L-loop graph, L momenta are undetermined and the Feynman rules state that one must integrate over all possible values of these L momenta.
- *Vertices:* the vertices describe the interactions among the particles in the theory. Vertices can in principle have any number of lines going in or out, but in many theories there are just cubic and/or quartic vertices. The Feynman rules translate each vertex into a mathematical rule; in the simplest case, this is just multiplication by a constant, but more generally the vertex rule can involve the momenta associated with the lines going in or out of the vertex.
- *Internal lines:* the Feynman rules translate every internal line into a "propagator" that depends on the momentum associated with the internal line. (A propagator is the Fourier transform of a Green's function.) The mathematical expression depends on the type of internal line, i.e. what kind of virtual particle is being exchanged.

One converts a Feynman diagram to a mathematical expression by tracing through the diagram and picking up factors, in successive order of appearance, from the rules for external and internal lines and vertices. The Feynman rules then state that $iA_n^{L\text{-loop}}$ is the sum of all possible L-loop diagrams with n external lines specified by the external data.

▷ *Example.* Consider the simplest case of a massless spin-0 scalar particle. The Feynman rules for external scalar lines are simply a factor of 1 while an internal line with momentum label P contributes $-i/P^2$; this is the massless scalar propagator. Let us assume that our scalars only interact via cubic vertices for which the Feynman rule is simply to multiply a factor of ig, where g is the coupling. A process involving four scalars then has the tree-level contribution

$$A_4^{\text{tree}}(1,2,3,4) =$$

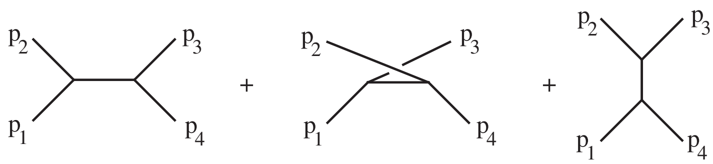

$$= g^2\left(\frac{1}{(p_1+p_2)^2} + \frac{1}{(p_1+p_3)^2} + \frac{1}{(p_1+p_4)^2}\right). \tag{1.10}$$

Here $(p_1+p_2)^2 = -(p_1^0+p_2^0)^2 + |\vec{p}_1+\vec{p}_2|^2$ etc. ◁

Note that the amplitude (1.10) is symmetric under exchange of identical bosonic external states. This is a manifestation of Bose symmetry. Similarly, Fermi statistics requires that an amplitude must be antisymmetric under exchange of any two identical external fermions.

▷ *Example.* A 1-loop diagram in our simple scalar model takes the form

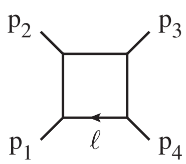

$$= g^4 \int \frac{d^4\ell}{(2\pi)^4} \frac{1}{\ell^2(\ell-p_1)^2(\ell-p_1-p_2)^2(\ell+p_4)^2}\,. \tag{1.11}$$

We have used momentum conservation $p_1+p_2+p_3+p_4=0$ to simplify the fourth propagator. The 1-loop 4-scalar amplitude in this model is obtained from the sum over inequivalent box diagrams (1.11). ◁

Loop-integrals can be divergent both in the large- and small-momentum regimes. Such "ultraviolet" and "infrared divergences" are well-understood and typically treated using a scheme such as dimensional regularization in which one replaces the measure $d^4\ell$ in the loop-integral by $d^D\ell$ with $D = 4-2\epsilon$. The divergences can then be cleanly extracted in the expansion of small ϵ. The treatment of ultraviolet divergences is called "renormalization"; it is a well-established method and the resulting predictions for scattering processes have been tested to high precision in particle physics experiments.

Example in QED

As an example of a tree-level process in QED, consider the annihilation of an electron–positron pair to two photons:

$$e^- + e^+ \;\to\; \gamma + \gamma$$

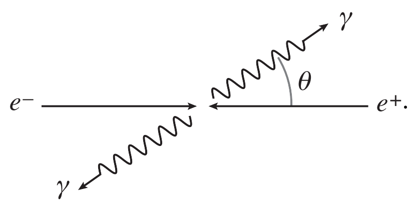

(1.12)

The leading-order contribution to the amplitude is a sum of two tree-level Feynman diagrams constructed from the basic QED vertex (1.8). It is

$$i A_4^{\text{tree}}\big(e^- + e^+ \;\to\; \gamma + \gamma\big) \;=$$

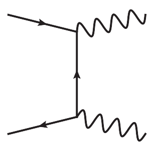

$+$

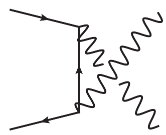

. (1.13)

The two diagrams (1.13) are translated via the QED Feynman rules to Lorentz-invariant contractions of 4-momentum vectors, photon polarizations, and fermion wavefunctions. In the standard formulation, the expression for the amplitude is not particularly illuminating. However, squaring it and subjecting it to a significant dose of index massage therapy, one finds that the scattering cross-section takes a rather simple form. In the high-energy limit, where the center of mass energy dominates the electron/positron rest mass, $E_{\text{CM}} \gg m_e c^2$, the result for the differential cross-section is

$$\frac{d\sigma}{d\Omega} \;=\; \frac{1}{64\pi^2 E_{\text{CM}}^2} \sum \big|A_4\big|^2 \;\xrightarrow{E_{\text{CM}} \gg m_e c^2}\; \alpha^2\, \frac{1}{E_{\text{CM}}^2}\, \frac{1 + \cos^2\theta}{1 - \cos^2\theta} + O(\alpha^4)\,. \tag{1.14}$$

The sum indicates a spin-sum average. The cross-section (1.14) depends on the center of mass energy E_{CM}, the scattering angle θ indicated in (1.12), and the fine structure constant α. (In natural units, $\alpha = e^2/4\pi$.) The result (1.14) serves to illustrate the salient properties of the differential cross-section: the dependence on particle energies and scattering angles as well as powers of a small fundamental dimensionless coupling constant α. In particular, the expression in (1.14) is the leading-order contribution to the scattering process, and higher orders, starting at 1-loop level, are indicated by $O(\alpha^4)$.

Since the focus in this book is on scattering amplitudes, and not the cross-sections, it is worthwhile to present an expression for the amplitude of the annihilation process $e^- + e^+ \to \gamma + \gamma$. It turns out to be particularly simple in the high-energy regime $E_{\text{CM}} \gg m_e c^2$, where the masses can be neglected and the momentum vectors of the incoming electron and positron can be chosen to be light-like (i.e. null, $p_i^2 = 0$). In that case, the sum of the two diagrams (1.13) reduces to a single term that can be written compactly as

$$A_4^{\text{tree}}\big(e^- + e^+ \to \gamma + \gamma\big) = 2e^2 \frac{\langle 24\rangle^2}{\langle 13\rangle\langle 23\rangle}\,. \tag{1.15}$$

Here, the angle brackets $\langle ij\rangle$ are closely related to the particle 4-momenta via $|\langle ij\rangle|^2 \sim 2p_i . p_j$, where the Lorentz-invariant dot-product is $p_i . p_j \equiv p_i^\mu p_{j\mu} = -p_i^0 p_j^0 + \vec{p}_i . \vec{p}_j$. The

4-momenta p_1 and p_2 are associated with the incoming electron and positron while p_3 and p_4 are the 4-momenta of the photons. The angle bracket notation $\langle ij\rangle$ is part of the ***spinor helicity formalism*** which is a powerful technical tool for describing scattering of massless particles in four spacetime dimensions. We will introduce the spinor helicity formalism in Chapter 2 and derive the expression (1.15) in full detail from the Feynman rules of QED. You may be surprised to note that the expression (1.15) is not symmetric in the momenta of the two photons. This is because the representation (1.15) selects distinct polarizations for the photons.

Using momentum conservation and the on-shell condition $p_i^2 = 0$, one can show that the norm-squared of the expression (1.15) is proportional to $(p_1.p_3)/(p_1.p_4)$. The sum in the cross-section (1.14) indicates a sum of the polarizations of the final state photons and an average over the spins of the initial state electron and positron. This procedure yields a second term with $p_3 \leftrightarrow p_4$. Hence, what goes into the formula (1.14) is

$$\sum \left|A_4^{\text{tree}}\big(e^- + e^+ \;\to\; \gamma + \gamma\big)\right|^2 = 2e^4\left(\frac{p_1.p_3}{p_1.p_4} + \frac{p_1.p_4}{p_1.p_3}\right). \tag{1.16}$$

An explicit representation of the 4-momenta is $p_1^\mu = (E, 0, 0, E)$ and $p_2^\mu = (E, 0, 0, -E)$ for the initial states and $p_3^\mu = (E, 0, E\sin\theta, E\cos\theta)$ and $p_4^\mu = (E, 0, -E\sin\theta, -E\cos\theta)$ for the final state photons. Momentum conservation is simply $p_1 + p_2 = p_3 + p_4$. Using that the center of mass energy is $E_{\text{CM}}^2 = -(p_1 + p_2)^2 = 4E^2$, one finds that the high-energy result for the differential cross-section in (1.14) follows from (1.16).

It is typical that amplitudes involving only massless particles are remarkably simpler than those with massive particles. This can be regarded as a high-energy regime, as in our example above. Because of the simplicity, yet remarkably rich mathematical structure, the focus of many developments in the field has been on scattering amplitudes of massless particles.

The S-matrix

We have introduced the scattering amplitude as a function that takes as its input the constrained external data $\{p_i^\mu, \text{type}_i\}$ and produces a complex number that is traditionally calculated in terms of Feynman diagrams. The scattering process can also be considered as an operation that maps an initial state $|i\rangle$ to a final state $|f\rangle$, each being a collection of single-particle states characterized by momenta and particle types. The scattering matrix S, also called the ***S-matrix***, is the unitary operator that "maps" the initial states to final states. In other words, the probability of an initial state $|i\rangle$ changing into a final state $|f\rangle$ is given by $|\langle f|S|i\rangle|^2$. Separating out the trivial part of the scattering process where no scattering occurs, we write

$$S = 1 + iT\,. \tag{1.17}$$

Then the amplitude is simply $A = \langle f|T|i\rangle$ and this is the quantity that is calculated by Feynman diagrams. Solving the S-matrix for a given theory means having a way to generate all scattering amplitudes at any order in perturbation theory.

Beyond Feynman diagrams

Armed with Feynman rules, we can, in principle, compute scattering amplitudes to our hearts' content. For instance, starting from the QED Lagrangian one can calculate the tree-level differential cross-section for the processes (1.1). It is typical for such a calculation that the starting point – the Lagrangian in its most compact form – is not too terribly complicated. And the final result for $d\sigma/d\Omega$ can be rather compact and simple too, as we have seen in (1.14). But the intermediate stages of the calculation often explode in an inferno of indices, contracted up-and-down and in all directions – providing little insight into the physics and hardly any hint of simplicity.

In a QFT course, students are exposed (hopefully!) to a lot of long character-building calculations, including the 4-particle QED processes in (1.1). But students are rarely asked to use standard Feynman rules to calculate processes that involve more than four or five particles, even at tree-level: for example, $e^- + e^+ \to e^- + e^+ + \gamma$ or $e^- + e^+ \to e^- + e^+ + \gamma + \gamma$. Why not? Well, one reason is that the number of Feynman diagrams tends to grow very quickly with the number of particles involved: for gluon scattering at *tree-level* in quantum chromodynamics (QCD) we have

$$\begin{array}{ll} g + g \to g + g & \text{4 diagrams} \\ g + g \to g + g + g & \text{25 diagrams} \\ g + g \to g + g + g + g & \text{220 diagrams} \end{array} \tag{1.18}$$

and for $g + g \to 8g$ one needs more than one million diagrams [1]. Another very important point is that the mathematical expression for each diagram becomes significantly more complicated as the number of external particles grows. So the reason students are not asked to calculate multi-gluon processes from Feynman diagrams is that it would be awful, un-insightful, and in many cases impossible.[2] There are tricks for simplifying the calculations; one is to use the spinor helicity formalism as indicated in our example for the high-energy limit of the tree-level process $e^- + e^+ \to \gamma + \gamma$. However, for a multi-gluon process even this does not directly provide a way to handle the growing number of increasingly complicated Feynman diagrams. Other methods are needed and you will learn much more about them in this book.

Thus, although visually appealing and seemingly intuitive, Feynman diagrams are not a particularly physical approach to studying amplitudes in gauge theories. The underlying reason is that the Feynman rules are non-unique: generally, individual Feynman diagrams are not physical observables. In theories, such as QED or QCD, that have gauge redundancy in their Lagrangian descriptions, we are required to fix the gauge in order to extract the Feynman rules from the Lagrangian. So the Feynman rules depend on the choice of gauge and therefore only gauge-invariant sums of diagrams are physically sensible. The on-shell amplitude is at each loop-order such a gauge-invariant sum of Feynman diagrams.

[2] Using computers to do the calculation can of course be very helpful, but not in all cases. Sometimes numerical evaluation of Feynman diagrams is simply so slow that it is not realistic to do. Moreover, given that there are poles that can cancel between diagrams, big numerical errors can arise in such evaluations. Therefore compact analytic expressions for the amplitudes are very useful in practical applications.

Furthermore, field redefinitions in the Lagrangian change the Feynman rules, but not the physics, hence the amplitudes are invariant. So field redefinitions and gauge choices can "move" the physical information between the diagrams in the Feynman expression for the amplitude and this can hugely obscure the appearance of its physical properties.

It turns out that despite the complications of the Feynman diagrams, the on-shell amplitudes for processes such as multi-gluon scattering $g + g \to g + g + \cdots + g$ can actually be written as remarkably simple expressions. This raises the questions: "why are the on-shell amplitudes so simple?" and "isn't there a better way to calculate amplitudes?". These are questions that have been explored in recent years and a lot of progress has been made towards improving calculational techniques and gaining insight into the underlying mathematical structure.

What do we mean by "mathematical structure" of amplitudes? At the simplest level, this means the analytic structure. For example, the amplitude (1.10) has a pole at $(p_1 + p_2)^2 = 0$ – this pole shows that a physical massless scalar particle is being exchanged in the process. Tree amplitudes are rational functions of the kinematic invariants, and understanding their pole structure is key to the derivation of on-shell recursive methods that provide a very efficient alternative to Feynman diagrams. These recursive methods allow one to construct on-shell n-particle tree amplitudes from input of on-shell amplitudes with fewer particles. The power of this approach is that gauge redundancy is eliminated since all input is manifestly on-shell and gauge invariant.

Loop amplitudes are integrals of rational functions and therefore have more complicated analytic structure, involving for example logarithmic branch cuts. The branch cuts can be exploited to reconstruct the loop amplitude from lower-loop and tree input, providing another class of efficient calculational techniques. On-shell approaches that exploit the analytic structure of amplitudes constitute a major theme in this book.

Thus, it is very useful to understand the analytic structure of amplitudes, but it is not the only mathematical structure of interest. We are also interested in the symmetries that leave the amplitudes invariant and how they constrain the result for the scattering process. Some amplitudes are invariant under symmetries that are not apparent in the Lagrangian and hence are not visible in the Feynman rules. Pursuing these symmetries and seeking to write the amplitudes in terms of variables that trivialize the constraints on the external data and simplify the action of the symmetries will often lead to new insights about the amplitudes; we will see this repeatedly as we develop the subject.

Finally, let us mention one of the highlights of what we mean by mathematical structure. It turns out that certain amplitudes naturally "live" in more abstract spaces in which they can be interpreted as volumes of geometric objects known as "polytopes"; these are higher-dimensional generalizations of polygons. The different compact representations found for a given amplitude turn out to correspond to different triangulations of the corresponding polytope. It is remarkable that a physical observable that encodes the probability for a scattering process as a function of energies and momenta has an alternative mathematical interpretation as the volume of a geometric object in a higher-dimensional abstract space! It is such surprising mathematical structures that are pursued in the third part of this book.

Some of the keywords for the topics we explore are:

1. spinor helicity formalism;
2. on-shell recursion relations (BCFW, CSW, all-line shifts, . . .);
3. on-shell superspace, superamplitudes, Ward identities;
4. twistors, zone-variables, momentum twistors;
5. dual superconformal symmetry and the Yangian;
6. generalized unitarity, maximal cuts;
7. Leading Singularities 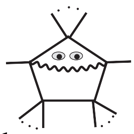 and on-shell blob-diagrams 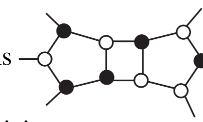;
8. the Grassmannian, polytopes, amplituhedrons, and mathematicians;
9. gravity $=$(gauge theory)^2, KLT and BCJ relations;

and much more.

The study of on-shell amplitudes may suggest a paradigm that can be phrased loosely as "avoiding the (full) Lagrangian" with all its ambiguities of field redefinitions and gauge choices, and instead focusing on how kinematics, symmetries, unitarity, and locality impact the physical observables. Or, more strongly, we may ask if the hints from the simplicity of amplitudes allow us to find another approach to perturbative quantum field theory: one might hope for a novel formulation that captures the physics of the full perturbative S-matrix. Such a new formulation could make amplitude calculations much more efficient and one might hope that it would lead to new insights even beyond amplitudes, for example for correlation functions of gauge-invariant operators and perhaps even for non-perturbative physics.

But now we are getting ahead of ourselves. The purpose of this book is to provide a practical introduction to on-shell methods for scattering amplitudes of massless particles. The choice of massless particles is because their processes are the simplest and currently the best understood. The above list of keywords will be explained and discussed in detail.

This book

Our presentation assumes a basic pre-knowledge of quantum field theory, in particular of the Dirac equation, simple scalar and QED Feynman rules and tree-level processes. Chapter 2 introduces the spinor helicity formalism in the context of tree-level Feynman rules and builds up the notation and conventions via explicit examples. Yang–Mills theory is introduced as are the needed tools for studying gluon amplitudes. Various fundamental properties of scattering amplitudes are discussed, including little group scaling and the analytic structure of tree amplitudes.

In Chapter 3 we take the first step towards modern on-shell methods for calculating scattering amplitudes, namely recursion relations for tree amplitudes. The material in Chapters 2 and 3 could be part of any modern course on quantum field theory, but traditionally it is not. So we hope that our presentation will be useful, either as part of a course or for self-study.