

The representation theory racah computes

A self-contained note built on the Killing–Cartan classification:
groups and global forms, irreps, Clebsch–Gordan coefficients,
recoupling, and gauge

The **racah** project

<https://github.com/Ryo-wtnb11/racah>

Notice: AI-generated document

This note was written by an AI agent as reference material for AI agents (and humans) working on the **racah** repository. **It may contain errors.** Check every statement against the code and the tests rather than trusting it blindly.

This note is *not* normative. It explains what the objects are and summarises the prior results the implementation rests on. The **frozen normative authority** for every basis, phase, ordering and normalisation convention — the actual numeric gauge — is [docs/gauge.md](#) (base $SU(2)$ and $SU(N)$) and [docs/gauge_soN.md](#) ($SO(N)/Sp(2N)$), verified by the golden tests. Where this note and those documents appear to disagree, those documents win. This note deliberately does not restate their numeric values.

Abstract

racah computes the coefficient data of the fusion category of finite-dimensional unitary representations of a connected compact simple Lie group. The note starts from the Killing–Cartan classification of the simple Lie algebras — $A_r, B_r, C_r, D_r, G_2, F_4, E_6, E_7, E_8$ — derives the compact groups and their global forms from it, and then defines the coefficients: irrep labels and dimensions, dualities and Frobenius–Schur indicators, fusion multiplicities, Clebsch–Gordan coefficients, and the recoupling (F) and braiding (R) symbols, for $SU(2)$, $SU(N)$, $Spin(N)/SO(N)$, $Sp(2N)$ and their central quotients. This note defines each of those objects, states the established results from the literature that the implementation uses, and records — as equations — the conventions the crate fixes where the mathematics leaves a choice. Scope is exactly what the implementation uses; it is not a course in Lie theory.

Contents

1	Scope, and how to read this note	2
2	Notation	3
3	The classification of the simple Lie algebras	3
3.1	Root systems, Cartan matrices, Dynkin diagrams	4
3.2	The nine diagrams, in Bourbaki numbering	4
3.3	Why the rank conventions start where they do	5
4	From algebras to compact groups, and to global forms	6
4.1	Step one: the compact simply connected group	6
4.2	Step two: global forms are centre quotients	6
4.3	Two consequences the implementation relies on	8

5	What <code>racah</code> implements, against the classification	8
6	Compact groups and their irreducible representations	9
7	Highest weights, Dynkin labels, dimensions	10
8	The families, their bases, and the two constructions	11
8.1	A -series: Gelfand–Tsetlin patterns	11
8.2	B, C, D : weights, spinors, and the generator bootstrap	12
8.3	$SU(2)$: closed forms	13
9	Duality and the Frobenius–Schur indicator	13
10	Tensor products, fusion multiplicities, outer multiplicity	14
11	Clebsch–Gordan coefficients	14
11.1	Orthonormality and completeness	15
11.2	Gauge freedom	15
12	Recoupling: $6j$, F -symbols, R -symbols	16
12.1	$SU(2)$: the $6j$ symbol	16
12.2	General G : the F -symbol as a four-CGC contraction	16
12.3	The R -symbol	17
13	The pentagon and hexagon identities	18
13.1	Pentagon	18
13.2	Hexagons	19
14	What is exact, and what is verification-gated	19

1 Scope, and how to read this note

`racah` is a coefficient library. Its outputs are numbers attached to tuples of irrep labels, and every one of those numbers is only defined *relative to a choice of basis*. The note is therefore organised around three questions, in order:

- (i) **What is the object?** A definition, with every symbol defined before use (Sections 3–13).
- (ii) **What does the literature already establish about it?** The prior results the implementation depends on, stated as facts with attribution. These are gathered in the section where they are used and flagged as **Fact**.
- (iii) **What does `racah` choose?** Where the mathematics leaves a free choice, the choice is stated here as an equation or a rule, and the normative document holding the frozen values is named.

The classification is the spine. The note is built on the Killing–Cartan classification, not on a list of groups. Section 3 states the classification and draws the nine Dynkin diagrams; Section 4 turns each family into a compact simply connected group with a known centre, and each subgroup of that centre into a global form; Section 5 then says which entries of the classification `racah` actually computes for and by which code path. Everything after that — highest weights, bases, fusion, CGC, F and R — refers back to those three sections rather than re-deriving where the families come from. This mirrors the crate: `group::RootSystem` carries all nine families and `group::GroupId` is a root system paired with a global form.

Two constructions, two normative documents. The crate has three coefficient engines but only two constructions. $SU(2)$ is evaluated from closed forms; $SU(N)$ is built on the Gelfand–Tsetlin basis; $SO(N)$ and $Sp(2N)$ are built by a generator bootstrap. Section 8 explains why the branching structure of each family forces the choice.

Citation numbering. Numbered citations $[n]$ in this note use the *same numbering* as [docs/references.md](#), which additionally carries the `file:symbol`-level port provenance. Entries $[1]$ – $[13]$ are that document’s bibliography verbatim; $[14]$ – $[21]$ are added here for results this note states that [docs/references.md](#) did not previously need to source. Every **Fact** and every mathematical table row carries a numbered citation with section-level precision; the tables of Section 5 are the one exception, and cite this repository’s own normative documents and tests instead, because they state facts about the crate rather than about mathematics.

2 Notation

Symbol	Meaning
G	a connected compact Lie group with simple Lie algebra
$\mathfrak{g}, r$	its Lie algebra (complexified where convenient), and the rank of $\mathfrak{g}$
$\mathfrak{h}, \mathfrak{h}^*$	a Cartan subalgebra and its dual (where weights live)
Φ, Φ^+, Δ	the roots, the positive roots, the simple roots $\alpha_1, \dots, \alpha_r$
$\alpha^\vee = 2\alpha/(\alpha, \alpha)$	the coroot of α
$\rho = \frac{1}{2} \sum_{\alpha \in \Phi^+} \alpha$	the Weyl vector
P, Q, P^+	the weight lattice, the root lattice, the dominant integral weights
W	the Weyl group
$\varepsilon_1, \dots, \varepsilon_n$	the orthonormal (ε) basis of $\mathfrak{h}^*$ used by the combinatorics
λ	a highest weight; $\lambda_i = (\lambda, \varepsilon_i)$ its ε -components
$a_i = \langle \lambda, \alpha_i^\vee \rangle$	the i -th Dynkin label of λ (Bourbaki numbering [20])
a, b, c, d, e, f	irreps of G , identified with their highest weights
$V_a, d_a = \dim V_a$	the carrier space of a and its dimension
$\bar{a}$	the dual (complex-conjugate) irrep of a
$\mathbf{1}$	the trivial irrep
m_a	a basis (“magnetic”) index inside V_a , $1 \leq m_a \leq d_a$
N_{ab}^c	the fusion multiplicity: how many copies of c occur in $a \otimes b$
$\mu, \nu, \kappa, \lambda$	outer-multiplicity indices; these are the four F -block multiplicity axes ¹
$C_{a m_a; b m_b}^{c \mu}$	a Clebsch–Gordan coefficient for the channel $(a, b \rightarrow c, \mu)$
P_{ab}^c	the orthogonal projector onto the c -isotypic component of $V_a \otimes V_b$
$[F_d^{abc}]_{(e, \mu \nu), (f, \kappa \lambda)}$	the F -symbol (associator) for fusing a, b, c into d
$[R_c^{ab}]_{\mu \nu}$	the R -symbol (braiding) exchanging a and b inside c
$\varkappa_a \in \{+1, -1, 0\}$	the Frobenius–Schur indicator of a
$Z(G), \Gamma$	the centre of G , and a subgroup of it that is quotiented out

Inner products $(\cdot, \cdot)$ on $\mathfrak{h}^*$ are the ones induced by the Killing form; $\langle \lambda, \alpha^\vee \rangle = 2(\lambda, \alpha)/(\alpha, \alpha)$. Overlines denote complex conjugation, daggers Hermitian adjoints.

3 The classification of the simple Lie algebras

Every group `racah` computes for is a compact form of one entry in a finite classification, and the crate’s type system is a transcription of that classification: `racah::group::RootSystem` has

¹The name λ is overloaded: it is a highest weight everywhere except on the F/R multiplicity axes, where it is the fourth multiplicity index. Context disambiguates, and $[\mu, \nu, \kappa, \lambda]$ is what the API calls those axes.

exactly nine variants, one per family. This section states the classification, draws its diagrams in the numbering the API uses, and records the rank conventions that make the list irredundant. It is the backbone the rest of the note hangs off.

3.1 Root systems, Cartan matrices, Dynkin diagrams

Let $\mathfrak{g}$ be a finite-dimensional complex Lie algebra. It is *simple* if it is non-abelian and has no proper non-zero ideal, and *semisimple* if it is a direct sum of simple ideals. Fix a Cartan subalgebra $\mathfrak{h} \subset \mathfrak{g}$ (a maximal self-normalising abelian subalgebra of semisimple elements); its dimension is the *rank* r . The adjoint action of $\mathfrak{h}$ decomposes $\mathfrak{g}$ into root spaces, giving the root system $\Phi \subset \mathfrak{h}^*$, and a choice of positive system fixes the simple roots $\Delta = \{\alpha_1, \dots, \alpha_r\}$ (Notation, Section 2). The *Cartan matrix* is

$$A_{ij} = \langle \alpha_j, \alpha_i^\vee \rangle = \frac{2(\alpha_i, \alpha_j)}{(\alpha_i, \alpha_i)} \in \mathbb{Z}, \quad A_{ii} = 2, \quad A_{ij} \leq 0 \ (i \neq j). \quad (1)$$

The *Dynkin diagram* of $\mathfrak{g}$ is the graph on r nodes with $A_{ij}A_{ji} \in \{0, 1, 2, 3\}$ edges between nodes $i \neq j$, decorated — when the two roots have different lengths — with an arrow pointing from the long root towards the short one. The diagram determines A , A determines $\mathfrak{g}$ up to isomorphism, and $\mathfrak{g}$ is simple exactly when its diagram is connected (Humphreys [8] §11.1–§11.2 for the Cartan matrix and the diagram, §14.2 for the isomorphism theorem; Bourbaki [20] Ch. VI §1.5, §4.1).

Theorem 3.1 (Killing–Cartan classification). *A complex simple Lie algebra is isomorphic to exactly one of*


$$A_r \ (r \geq 1), \quad B_r \ (r \geq 2), \quad C_r \ (r \geq 3), \quad D_r \ (r \geq 4), \quad G_2, \quad F_4, \quad E_6, \quad E_7, \quad E_8,$$

with the Dynkin diagrams of Section 3.2. Equivalently: the connected Dynkin diagrams are exactly those nine shapes. Humphreys [8] §11.4 (Theorem) with the proof of §11.4–§11.5 and the existence half in §12.1; Bourbaki [20] Ch. VI §4.2 (Theorem 3) and Plates I–IX; Fulton–Harris [7] §21.1.


The first four are the *classical* families — they are the ones with a uniform matrix description, and the only ones `racah` computes coefficients for — and the last five are the *exceptional* algebras.

3.2 The nine diagrams, in Bourbaki numbering


Nodes are drawn $\circ$ and labelled underneath with their Bourbaki index; the index order is exactly the order of the Dynkin label tuple $(a_1, \dots, a_r)$ that the API accepts and returns (Convention 7.2). A double or triple bond carries an arrow pointing at the *short* root.




$D_r, r \geq 4$ — a chain of $r - 2$ nodes forking into two:



E_8 — a chain of seven nodes with node 2 attached to node 4:



E_7 is the subdiagram on nodes $\{1, \dots, 7\}$ and E_6 the subdiagram on $\{1, \dots, 6\}$ — deleting node 8, then node 7, from the tail of the chain (Bourbaki [20] Plates V–VII):



3.3 Why the rank conventions start where they do

The rank bounds in Theorem 3.1 are not arbitrary: at low rank the diagram shapes coincide, and without the bounds the list would repeat itself. The coincidences are exactly

Diagram coincidence	Algebra isomorphism	Group statement
$A_1 = B_1 = C_1$	$\mathfrak{su}(2) \cong \mathfrak{so}(3) \cong \mathfrak{sp}(2)$	$SU(2) \cong Spin(3) \cong Sp(2)$
$B_2 = C_2$	$\mathfrak{so}(5) \cong \mathfrak{sp}(4)$	$Spin(5) \cong Sp(4)$
$D_2 = A_1 \times A_1$	$\mathfrak{so}(4) \cong \mathfrak{su}(2) \oplus \mathfrak{su}(2)$	$Spin(4) \cong SU(2) \times SU(2)$
$D_3 = A_3$	$\mathfrak{so}(6) \cong \mathfrak{su}(4)$	$Spin(6) \cong SU(4)$

D_2 is the only one of the four that is not simple at all: its diagram is disconnected, so $\mathfrak{so}(4)$ is semisimple but not simple, and the theorem does not apply to it. Each line is a diagram identity first — read off the shapes above at $r = 1, 2, 3$ — and the algebra and group statements follow from Theorem 3.1 and Fact 4.1. Bourbaki [20] Ch. VI §4.4 and the Plates I–IV remarks (the identifications of low-rank systems); Fulton–Harris [7] §18 (the classical algebras, where the same isomorphisms are exhibited on the defining representations).

Remark 3.2 (What the code does with the coincidences). `racah` takes the coincidences seriously in the $B/C/D$ bootstrap but does *not* adopt Bourbaki’s irredundant bounds verbatim, because a constructor is named after a *group*, not a diagram. The bootstrap family key `bcd::Series` enforces $r \geq 2$ for B and C and $r \geq 3$ for D : the genuinely duplicated low ranks B_1 , C_1 , D_2 are rejected with a per-global-form redirect ($Spin(3) \cong SU(2)$, $Sp(2) \cong SU(2)$, $Spin(4) \cong SU(2) \times SU(2)$, and for the quotients $SO(3)/PSp(2) = \text{integer } j \text{ only}$, $SO(4) = j_1 + j_2 \text{ integer}$), while C_2 and D_3 are *kept* — $Sp(4)$ and $SO(6)$ are groups a caller legitimately asks for by name, and the bootstrap builds them directly rather than routing the caller through $Spin(5)$ or $SU(4)$. Nothing downstream depends on the redundancy being removed: the two presentations of the same group give the same coefficients up to the relabelling, and only the label conventions differ. API: `bcd::Series::min_rank`, `bcd::Series::low_rank_redirect`.

Remark 3.3 (Rank versus matrix dimension). Two integers are in play and they are not the same. r is the rank — the number of Dynkin labels. N is the matrix dimension of the defining representation. The `GroupId` constructors take N (`GroupId::so(7)`, `GroupId::sp(6)`) and convert: $SU(N) \leftrightarrow A_{N-1}$, $Spin(2r+1) \leftrightarrow B_r$, $Sp(2r) \leftrightarrow C_r$, $Spin(2r) \leftrightarrow D_r$. Weight tuples are always indexed by r .

4 From algebras to compact groups, and to global forms

Theorem 3.1 classifies complex *algebras*. `racah` computes for *compact groups*, and the passage from one to the other is two steps, neither of them unique in the naive sense: first a compact real form and its simply connected cover, then a quotient by a subgroup of the centre.

4.1 Step one: the compact simply connected group

Fact 4.1 (Compact real forms and their covers). Every complex simple Lie algebra $\mathfrak{g}$ has a compact real form $\mathfrak{k} \subset \mathfrak{g}$, unique up to conjugation, and there is a unique connected simply connected compact Lie group G_{sc} with Lie algebra $\mathfrak{k}$. Its centre is finite, and

$$Z(G_{\text{sc}}) \cong P/Q, \quad (2)$$

the weight lattice modulo the root lattice — a finite abelian group read straight off the Cartan matrix (1), since Q is spanned by the simple roots and P by the fundamental weights. Existence and conjugacy-uniqueness of the compact real form: Knapp [21] Theorem 6.11 and Corollary 6.20. The centre: Bröcker–tom Dieck [14] §V.7; Humphreys [8] §13.1 for the P/Q computation itself.

Applying (2) family by family gives the table this note and the crate both rest on. It is the dictionary between the classification and the group names in the API.

Family	Compact simply connected G_{sc}	$Z(G_{\text{sc}}) \cong P/Q$	Adjoint form G_{sc}/Z
$A_{N-1}, N \geq 2$	$SU(N)$	$\mathbb{Z}_N$	$PSU(N)$
$B_r, r \geq 2$	$Spin(2r+1)$	$\mathbb{Z}_2$	$SO(2r+1) = PSO(2r+1)$
$C_r, r \geq 2$	$Sp(2r)$	$\mathbb{Z}_2$	$PSp(2r)$
D_r, r odd	$Spin(2r)$	$\mathbb{Z}_4$	$PSO(2r)$
D_r, r even	$Spin(2r)$	$\mathbb{Z}_2 \times \mathbb{Z}_2$	$PSO(2r)$
G_2	G_2	1	G_2
F_4	F_4	1	F_4
E_6	E_6	$\mathbb{Z}_3$	$E_6/\mathbb{Z}_3$
E_7	E_7	$\mathbb{Z}_2$	$E_7/\mathbb{Z}_2$
E_8	E_8	1	E_8

Sources for the table. Each P/Q is tabulated system by system in Bourbaki [20] Plates I–IX (the plate for each family lists P/Q and the fundamental weights); the identification with $Z(G_{\text{sc}})$ is Fact 4.1. The D_r split is the arithmetic of the plate for D_r : P/Q is generated by $[\omega_1], [\omega_{r-1}], [\omega_r]$ with $2[\omega_r] = [\omega_1]$ when r is odd (so the group is cyclic of order 4) and $2[\omega_r] = 0$ when r is even (so it is $\mathbb{Z}_2 \times \mathbb{Z}_2$) — the parity dependence traces to $(\omega_r, \omega_r) = r/4$. The congruency classes of every family are also tabulated by Slansky [13] §5.

Three families (G_2, F_4, E_8) have trivial centre, so for them “simply connected” and “adjoint” name the same group and there is no global-form question to ask at all. These centres are recorded in the `RootSystem` variant docs and are the reason the exceptional variants carry no admissibility logic (Section 5).

4.2 Step two: global forms are centre quotients

Everything the rest of this note computes is stated for a group, but a root system names only an algebra. $\mathfrak{so}(N)$ is the algebra of both $Spin(N)$ and $SO(N) = Spin(N)/\mathbb{Z}_2$; $\mathfrak{su}(N)$ is the algebra of $SU(N)$, of $SU(N)/\mathbb{Z}_k$ for every $k \mid N$, and of $PSU(N) = SU(N)/\mathbb{Z}_N$. These groups have the same Lie-algebra representation theory and differ in exactly one respect: which dominant integral weights are genuine representations *of the group*.

Fact 4.2 (Classification of connected compact forms). Let G_{sc} be the simply connected compact group with a given root system. Then $Z(G_{\text{sc}}) \cong P/Q$, every connected compact group with that root system is $G_\Gamma = G_{\text{sc}}/\Gamma$ for a subgroup $\Gamma \subseteq Z(G_{\text{sc}})$, and

$$\text{Irr}(G_\Gamma) = \{ \lambda \in P^+ : \chi_\lambda|_\Gamma = 1 \}, \quad (3)$$

where the *central character* χ_λ depends on λ only through its class $[\lambda] \in P/Q$. [7] §23; the congruency classes are tabulated by Slansky [13] §5.

The class $[\lambda] \in P/Q$ is the *N-ality* / congruency class, and for the series **racah** implements it is an explicit linear form in the Dynkin labels.

$A_r, r = N - 1$: $P/Q \cong \mathbb{Z}_N$. The class is the *N-ality*

$$t(\lambda) \equiv \sum_{i=1}^{N-1} i a_i \pmod{N}, \quad (4)$$

and for $\Gamma = \mathbb{Z}_k$ with $k \mid N$ the admissibility condition (3) reads $t(\lambda) \equiv 0 \pmod{k}$. So $PSU(3) = SU(3)/\mathbb{Z}_3$ admits the adjoint **8** and rejects the fundamental **3**.

B_r : $P/Q \cong \mathbb{Z}_2$. The class is $a_r \pmod{2}$ — exactly the tensor/spinor split described in Section 8.2. $Spin(2r + 1)$ admits all $\lambda \in P^+$; $SO(2r + 1)$ admits those with a_r even.

C_r : $P/Q \cong \mathbb{Z}_2$. The class is $a_1 + a_3 + a_5 + \dots \pmod{2}$. $Sp(2r)$ admits all $\lambda \in P^+$; $PSp(2r) = Sp(2r)/\mathbb{Z}_2$ admits those with $a_1 + a_3 + \dots$ even.

D_r : $P/Q \cong \mathbb{Z}_4$ (*r odd*) or $\mathbb{Z}_2 \times \mathbb{Z}_2$ (*r even*). This is the rich case, and it is where a convention has to be pinned rather than derived, because naming a form requires choosing a generator of the centre.

One condition is uniform in r : the *vector* (tensor) sublattice $\{0, v\} \subset P/Q$ is always a $\mathbb{Z}_2$, and quotienting by it gives

$$SO(2r) : \quad a_{r-1} + a_r \equiv 0 \pmod{2}, \quad (5)$$

for every r — this is the tensor/spinor split described in Section 8.2, and it is r -uniform because $(\varepsilon_1, \varepsilon_1) = 1$ for every r . The remaining forms depend on the parity of r .

For r odd the class is valued in $\mathbb{Z}_4$, and **racah** fixes the generator $c = [\omega_r]$, so that

$$\kappa(\lambda) \equiv a_r - a_{r-1} + 2 \sum_{\substack{i \text{ odd} \\ i \leq r-2}} a_i \pmod{4}. \quad (6)$$

This choice makes conjugation *be* negation of the class, matching **bcd::Irrep::dual** ($\lambda_r \mapsto -\lambda_r$). Slansky's uniform formula [13] §5 equals $+\kappa$ for $r \equiv 1 \pmod{4}$ and $-\kappa$ for $r \equiv 3 \pmod{4}$; the difference is the $\mathbb{Z}_4$ automorphism $x \mapsto -x$, i.e. a generator choice, not a discrepancy.

so (5) is the $\kappa \equiv 0 \pmod{2}$ subgroup, and $PSO(2r)$ (the full $\mathbb{Z}_4$ quotient) is $\kappa \equiv 0 \pmod{4}$. There are no half-spin forms for r odd: $\mathbb{Z}_4$ has only one subgroup of order 2.

For r even the class is valued in $\mathbb{Z}_2 \times \mathbb{Z}_2$, which has three subgroups of order 2, so alongside $Spin(2r)$ (all λ), $SO(2r)$ (5) and $PSO(2r)$ (the full quotient) there are two *half-spin* forms. Writing $t \equiv \sum_{i \text{ odd}, i \leq r-2} a_i \pmod{2}$ and $p \equiv a_{r-1} + t$, $q \equiv a_r + t \pmod{2}$:

$$\text{half-spin}_+ : p \equiv 0; \quad \text{half-spin}_- : q \equiv 0; \quad PSO(2r) : p \equiv q \equiv 0 \pmod{2}.$$

A half-spin form admits one chirality of spinor but *not* the vector representation (which has $p = q = 1$). The two half-spin forms are named for the class they *retain*, not the central element they kill: a kill-based name is not uniform in r , because $(\omega_r, \omega_r) = r/4$ makes the annihilator of z_{ω_r} differ between $r \equiv 0$ and $r \equiv 2 \pmod{4}$. Slansky's S_s/S_c are accepted as documented aliases. [docs/references.md](#) carries this convention with its full justification.

4.3 Two consequences the implementation relies on

Proposition 4.3 (Admissibility is a construction-time predicate). *The class map $P^+ \rightarrow P/Q$, $\lambda \mapsto [\lambda]$, is the restriction of a group homomorphism, so $[\lambda + \mu] = [\lambda] + [\mu]$ and $[-w_0\lambda] = -[\lambda]$. Hence the admissible set of (3) is closed under fusion and under duality: it never needs re-checking on a fusion output.*

Proposition 4.4 (A global form deletes irreps; it never changes a coefficient). *Let λ be admissible for both G_Γ and $G_{\Gamma'}$. The representations of G_Γ and $G_{\Gamma'}$ with highest weight λ have the same carrier, the same $\mathfrak{g}$ -action, and hence the same CGC, F , R and $\varkappa$ in the same basis: the choice of Γ never enters their computation. Consequently the coefficient caches are keyed on irrep labels with no global form in the key, and two forms sharing an irrep must share the cache entry.*

API: `racah::group` — `RootSystem`, `GlobalForm`, `CenterSubgroup`, `GroupId`, and the predicate `GroupId::admits`; the form-aware constructors `sun::Irrep::from_dynkin_in` and `bcd::Irrep::from_dynkin_in` (the plain `from_dynkin` constructors keep their historically published groups).

Remark 4.5 (Scope). Connected compact groups only. $O(N)$ and $Pin(N)$ are not central quotients of a simply connected group — they are disconnected extensions — and are out of scope; so is any non-compact real form.

5 What racah implements, against the classification

Sections 3 and 4 describe every simple compact group. `racah` does not compute for all of them. This section is the correspondence table — which entry of the classification is served by which code path, and what happens at the entries that are not served. It cites the repository’s own normative documents and tests rather than the literature: these are facts about this crate, not about mathematics.

Family	Groups	Engine	Where it lives
A_1	$SU(2), SO(3)$	closed forms (Section 8.3)	<code>su2</code> ; exact $3j/6j$; docs/gauge.md
A_{N-1}	$SU(N)$, $SU(N)/\mathbb{Z}_k$, $PSU(N)$	Gelfand–Tsetlin (Section 8.1)	<code>sun</code> ; gauge docs/gauge.md
B_r	$Spin(2r+1)$, $SO(2r+1)$	generator bootstrap (Section 8.2)	<code>bcd</code> , <code>Series::B</code> ; docs/gauge_soN.md
C_r	$Sp(2r), PSp(2r)$	generator bootstrap	<code>bcd</code> , <code>Series::C</code> ; docs/gauge_soN.md
D_r	$Spin(2r)$, $SO(2r)$, $PSO(2r)$, half-spin	generator bootstrap	<code>bcd</code> , <code>Series::D</code> ; docs/gauge_soN.md
G_2, F_4, E_6, E_7, E_8	—	none	<code>RootSystem</code> variants only

Two types, two jobs. `RootSystem` is the *label-lattice* type: it carries all nine families of Theorem 3.1, it fixes the rank (hence the length of the Dynkin tuple) and the lattice P/Q (hence the congruences of Section 4). `bcd::Series` is the *bootstrap-family* key: it exists only where a defining-representation seed and a canonical catalogue exist, so it has exactly three variants B, C, D . The single bridge is `bcd::Series::root_system`; neither type is an alias of the other. A -series labels are not a `Series` at all — they go through `sun::Irrep` (and A_1 through `su2`).

A group is a root system plus a global form. The whole of `GroupId` is

$$\begin{aligned} \text{GroupId} &= \text{RootSystem} \times \text{GlobalForm}, \\ \text{GlobalForm} &\in \{ \text{SimplyConnected} \} \cup \{ \text{Quotient}(\Gamma) : \Gamma \subseteq Z \}, \end{aligned}$$

which is exactly Fact 4.2: the classification supplies the first factor, the centre supplies the second. The named constructors are the ergonomic surface, and they take the matrix dimension, not the rank (all are associated functions of `GroupId`, and the prefix is dropped in the table):

Constructor	Group	Root datum
<code>su(n)</code>	$SU(n)$	A_{n-1} , <code>SimplyConnected</code>
<code>su_quotient(n,k)</code>	$SU(n)/\mathbb{Z}_k$, $k \mid n$	A_{n-1} , $\mathbb{Z}k(k)$
<code>psu(n)</code>	$PSU(n)$	A_{n-1} , $\mathbb{Z}k(n)$
<code>spin(n)</code>	$Spin(n)$	$B_{(n-1)/2}$ or $D_{n/2}$, <code>SimplyConnected</code>
<code>so(n)</code>	$SO(n)$	$B_{(n-1)/2} + \mathbb{Z}2$, or $D_{n/2} + \text{DVector}$
<code>pso(n)</code>	$PSO(n)$	$B + \mathbb{Z}2$; $D_r + \mathbb{Z}4$ (r odd), <code>DFull</code> (r even)
<code>sp(m)</code>	$Sp(m)$	$C_{m/2}$, <code>SimplyConnected</code>
<code>psp(m)</code>	$PSp(m)$	$C_{m/2}$, $\mathbb{Z}2$
<code>half_spin_plus(n)</code>	$Ss(n)$, $4 \mid n$	$D_{n/2}$, <code>DHalfSpinPlus</code>
<code>half_spin_minus(n)</code>	$Sc(n)$, $4 \mid n$	$D_{n/2}$, <code>DHalfSpinMinus</code>

Which form the coefficient engines publish. A construction engine works in the *cover* — the bootstrap needs the spinor seeds, so its internal catalogue is over $Spin(2r+1)$, $Sp(2r)$, $Spin(2r)$ — while the labels a query *publishes* are cut down to a chosen default form: $SO(2r+1)$ for B , $Sp(2r)$ for C , $SO(2r)$ for D . The form-aware constructors `sun::Irrep::from_dynkin_in` and `bcd::Irrep::from_dynkin_in` take a `GroupId` explicitly; the plain `from_dynkin` constructors keep those historically published defaults. API: `bcd::Series::cover_group`, `bcd::Series::published_group`.

The exceptionals are placeholders, and provably inert. G_2 , F_4 , E_6 , E_7 , E_8 exist as `RootSystem` variants so that the type transcribes Theorem 3.1 completely — rank and centre are recorded, nothing else. There is no coefficient engine and no congruence logic: `GroupId::admits` answers `true` for `SimplyConnected` (the cover admits every dominant integral weight, which needs no family-specific arithmetic) and `false` for *every* quotient, including the mathematically real ones such as $E_6/\mathbb{Z}_3$. That is not a silent wrong answer but the deliberate fail-closed behaviour of Section 4: a group with no admissible weights has no irreps, so both form-aware constructors reject every label with the ordinary not-admissible error. It is pinned by the test `tests/global_form.rs::nonsensical_root_datum_admits_nothing`, which sweeps labels for a list of impossible root data — E_6 with $\mathbb{Z}_3$ among them — and asserts that nothing is admitted.

6 Compact groups and their irreducible representations

Let G be a compact Lie group and V a finite-dimensional complex vector space. A *representation* is a continuous homomorphism $\pi : G \rightarrow GL(V)$; it is *unitary* if $\pi(g)^\dagger \pi(g) = \mathbb{I}$ for all g , and *irreducible* (an *irrep*) if the only π -invariant subspaces of V are 0 and V . Write $\text{Irr}(G)$ for the set of isomorphism classes of irreps.

The whole subject rests on three classical results, which are what make “coefficient library” a well-posed idea at all.

Fact 6.1 (Unitarisability and complete reducibility; Weyl). Every finite-dimensional representation of a compact group admits an invariant Hermitian inner product (average any inner product against the Haar measure), and decomposes as a direct sum of irreps. All irreps are finite-dimensional. See Bröcker–tom Dieck [14] §II.1–II.4, or Fulton–Harris [7] §9.

Consequently every representation **racah** touches may be taken unitary with an orthonormal basis, and every tensor product decomposes.

Fact 6.2 (Schur’s lemma). If a, b are irreps then $\text{Hom}_G(V_a, V_b)$ — the space of linear maps commuting with the action — is $\mathbb{C}$ if $a \cong b$ and 0 otherwise. More generally, for representations V, W , $\dim \text{Hom}_G(V, W)$ counts the common irreducible constituents with multiplicity. [14] §II.1, [7] §1.2.

Fact 6.3 (Peter–Weyl and Schur orthogonality). The matrix coefficients $\sqrt{d_a} \pi_{mn}^a$ over all $a \in \text{Irr}(G)$ form an orthonormal basis of $L^2(G)$ with respect to normalised Haar measure; in particular

$$\int_G \overline{\pi_{mn}^a(g)} \pi_{m'n'}^b(g) dg = \frac{1}{d_a} \delta_{ab} \delta_{mm'} \delta_{nn'}.$$

[14] §III.3, [7] §9. Schur orthogonality is the source of the Clebsch–Gordan orthonormality and completeness relations of Section 11: those relations are not conventions, they are theorems.

Restricting attention to G connected with $\mathfrak{g}$ simple, $\text{Irr}(G)$ is classified by highest weights, which is the subject of the next section.

7 Highest weights, Dynkin labels, dimensions

Fix a maximal torus $T \subset G$ with Lie algebra $\mathfrak{t}$, and let $\mathfrak{h} = \mathfrak{t} \otimes \mathbb{C}$ be the corresponding Cartan subalgebra of $\mathfrak{g}_{\mathbb{C}}$. A representation V decomposes into simultaneous eigenspaces of $\mathfrak{h}$,

$$V = \bigoplus_{\omega \in \mathfrak{h}^*} V_{\omega}, \quad V_{\omega} = \{v \in V : H \cdot v = \omega(H)v \ \forall H \in \mathfrak{h}\},$$

and the ω with $V_{\omega} \neq 0$ are the *weights* of V , with *multiplicity* $\dim V_{\omega}$. The weights of the adjoint representation other than 0 are the *roots* Φ . A choice of positive system Φ^+ picks out simple roots $\Delta = \{\alpha_1, \dots, \alpha_r\}$; a weight λ is *dominant integral* if $a_i := \langle \lambda, \alpha_i^{\vee} \rangle \in \mathbb{Z}_{\geq 0}$ for all i , and the tuple $(a_1, \dots, a_r)$ is its vector of *Dynkin labels*.

Fact 7.1 (Theorem of the highest weight; Cartan–Weyl). For G connected compact simply connected with simple $\mathfrak{g}$, the map $V \mapsto$ (its highest weight, in the dominance order induced by Φ^+) is a bijection $\text{Irr}(G) \rightarrow P^+$. Each irrep has a one-dimensional highest weight space, annihilated by every raising operator e_{α_i} . Humphreys [8] §20–21; [7] §14.

Convention 7.2 (Labelling). **racah** accepts and returns irreps as Dynkin label tuples $(a_1, \dots, a_r)$ in **Bourbaki numbering** [20] (so, e.g., for B_r the short simple root is α_r and a_r is the spinor label). Internally the combinatorics use the ε -basis components $\lambda = (\lambda_1 \geq \lambda_2 \geq \dots)$; the two are related by the series-specific linear map, and both directions are exact integer (or half-integer, for spinors) arithmetic. API: `sun::Irrep::from_dynkin`, `bcd::Irrep::from_dynkin`.

Fact 7.3 (Weyl dimension formula). For $\lambda \in P^+$,

$$d_{\lambda} = \prod_{\alpha \in \Phi^+} \frac{(\lambda + \rho, \alpha)}{(\rho, \alpha)}. \quad (7)$$

[8] §24.3, [7] §24.3. The right-hand side is a ratio of integers and is evaluated exactly.

Fact 7.4 (Freudenthal’s recursion). The weight multiplicities $m_{\lambda}(\omega) = \dim V_{\omega}$ of V_{λ} satisfy

$$(\|\lambda + \rho\|^2 - \|\omega + \rho\|^2) m_{\lambda}(\omega) = 2 \sum_{\alpha \in \Phi^+} \sum_{k \geq 1} m_{\lambda}(\omega + k\alpha) (\omega + k\alpha, \alpha), \quad (8)$$

which determines every $m_{\lambda}(\omega)$ downward from $m_{\lambda}(\lambda) = 1$. [8] §22.3.

Equations (7) and (8) are pure integer combinatorics; **racah** computes both in exact arithmetic, with QSpace [3] used as the independent numerical dimension oracle for $B/C/D$.

8 The families, their bases, and the two constructions

The coefficients of Sections 11–13 are defined only once an orthonormal basis has been chosen inside every V_a . The two constructions `racah` implements are two ways of producing that basis, and which one a family gets is forced by its branching structure.

8.1 A-series: Gelfand–Tsetlin patterns

Fact 8.1 (Weyl branching for the unitary chain). Restricting an irrep of $U(k)$ with highest weight $\lambda = (\lambda_1 \geq \dots \geq \lambda_k)$ to $U(k-1)$ gives

$$V_\lambda^{U(k)}|_{U(k-1)} = \bigoplus_{\mu} V_\mu^{U(k-1)}, \quad \lambda_1 \geq \mu_1 \geq \lambda_2 \geq \mu_2 \geq \dots \geq \mu_{k-1} \geq \lambda_k,$$

each admissible μ occurring *exactly once*: the chain is *multiplicity-free*. [7] §6.1, §15.3.

Iterating $U(N) \supset U(N-1) \supset \dots \supset U(1)$ therefore labels each basis vector of an $SU(N)$ irrep uniquely by the whole tower of intermediate highest weights — a triangular array

$$m = (m_{k,l})_{1 \leq k \leq l \leq N}, \quad m_{k,l+1} \geq m_{k,l} \geq m_{k+1,l+1},$$

the *Gelfand–Tsetlin (GT) pattern*, whose top row $l = N$ is λ .

Fact 8.2 (Gelfand–Tsetlin). The GT patterns with top row λ index an orthonormal basis of V_λ , and the generators of $\mathfrak{gl}_N$ act on it by explicit closed-form matrix elements: the Cartan generators are diagonal with

$$\text{wt}(m)_l = \text{rowsum}(l) - \text{rowsum}(l-1), \quad \text{rowsum}(l) = \sum_{k=1}^l m_{k,l},$$

and the elementary raising/lowering operators $E_{l,l+1}$, $E_{l+1,l}$ have matrix elements given by a ratio of products of linear forms in the pattern entries. Original statement: Gelfand–Tsetlin [4]; the explicit matrix elements and their derivation: Baird–Biedenharn [15]; a modern account covering the classical series: Molev [12].

The uniqueness of the labelling is what makes the ladder elements closed-form, and that closed form is what makes a direct, exact CGC construction possible. Two points are worth stating because they are frequently confused:

Remark 8.3. The SU chain *alone* is not multiplicity-free: the adjoint **8** of $SU(3)$ restricts to $SU(2)$ as $\mathbf{2} \oplus \mathbf{2} \oplus \mathbf{3} \oplus \mathbf{1}$. It is the intermediate $U(1)$ charge at each step that separates the recurring copies, which is why the chain is written with $U(k)$, not $SU(k)$.

Convention 8.4 (GT basis order). The magnetic indices m_a of every $SU(N)$ coefficient are positions in `racah`’s GT enumeration order, which reproduces `SUNRepresentations.jl` [9] index-for-index and is pinned by checked-in fixtures. This order is *load-bearing*: a re-enumeration emitting the same set of patterns in a different order permutes every returned coefficient. The normative statement, including which pattern occupies the last basis index, is [docs/gauge.md](#) §1.

The construction is then: enumerate GT patterns; build the exact rational ladder matrices of Fact 8.2; solve the highest-weight nullspace; gauge-fix; descend by the lowering operators. This is the algorithm of Alex, Kalus, Huckleberry and von Delft [1], whose reference implementation `SUNRepresentations.jl` [9] `racah` ports. Steps and tolerances: [docs/gauge.md](#) §2–§7.

8.2 B, C, D : weights, spinors, and the generator bootstrap

For the orthogonal and symplectic series the ε -basis description is the practical one. With $\lambda = \sum_i \lambda_i \varepsilon_i$:

- $B_r = \mathfrak{so}(2r+1)$: $\lambda_1 \geq \dots \geq \lambda_r \geq 0$, all λ_i simultaneously integers (*tensor* irreps) or all half-odd-integers (*spinor* irreps). The Dynkin label a_r is even in the tensor sector, odd in the spinor sector.
- $C_r = \mathfrak{sp}(2r)$: $\lambda_1 \geq \dots \geq \lambda_r \geq 0$, all integers. There is no spinor sector; the $\mathbb{Z}_2$ centre is detected by $a_1 + a_3 + a_5 + \dots$.
- $D_r = \mathfrak{so}(2r)$: $\lambda_1 \geq \dots \geq \lambda_{r-1} \geq |\lambda_r|$, with λ_r of *either sign* — the sign distinguishes the two chiralities, and swapping it is the outer automorphism exchanging the two half-spin fundamentals ω_{r-1} and ω_r . Again all λ_i integral (tensor) or all half-odd-integral (spinor).

Root data, the $\varepsilon \leftrightarrow$ Dynkin maps, and the spinor sectors are Fulton–Harris [7] §§18–20, 24 and Humphreys [8] §12–13. Low ranks where the classification degenerates (B_1, C_1, D_2) are rejected and redirected by the API rather than silently aliased.

Why not GT here.

Fact 8.5 (Branching for the orthogonal and symplectic chains). $SO(n) \supset SO(n-1)$ is multiplicity-free, and explicit GT-type bases with closed-form matrix elements exist for it [4, 12]. The symplectic chain $Sp(2r) \supset Sp(2r-2)$ is *not* multiplicity-free: intermediate irreps recur, so no GT-type pattern labels states uniquely there, and no practical closed-form ladder matrix elements follow. Molev [12] §§2–4 gives both statements and the (substantially more involved) orthogonal formulas.

So a GT route would cover B and D only, at high cost and with no production reference implementation, and would still leave C unserved. `racah` therefore uses a single *generator bootstrap* for all of B, C, D , following the discipline of Weichselbaum’s QSpace [2, 3]:

1. **Seed** the defining representation of each series explicitly — simple-root raising operators and Cartan generators, written out per series — plus, for B and D , the Clifford-algebra seeds of the fundamental spinors, which lie in no tensor power of the defining representation and so must be base cases.
2. **Multiply**: form tensor products of already-materialised irreps.
3. **Decompose** numerically: sweep for highest-weight vectors (a nullspace/Gram–Schmidt sweep), orthonormalise by a positive-diagonal QR.
4. **Harvest** the new irreps’ generators and **recurse**.

The price of that generality is that the basis — and hence the gauge — is defined *procedurally* rather than by a formula: it is whatever the deterministic sweep produces. [docs/gauge_soN.md](#) pins that determinism down, including the descending-weight sort and its tie-break (§7), the first-significant-entry block sign convention (§8), the outer-multiplicity assignment in discovery order (§9), and the well-order $\prec$ on irreps together with the *canonical-parent rule* (§14) that makes “which product an irrep is harvested from” a function of the label rather than of the call history.

Convention 8.6 (Canonical parent). An irrep’s generators are materialised from one distinguished product, selected by a total order $\prec$ on irreps of a fixed (series, rank) that compares doubled ε -basis box count first, then exact Weyl dimension, then the Dynkin tuple lexicographically. $\prec$ is a well-order, which is what makes the on-demand recursion terminate. The normative statement, and the reason box count rather than dimension is the primary key, is [docs/gauge_soN.md §14.1–14.2](#).

8.3 $SU(2)$: closed forms

For $SU(2)$ nothing is constructed. The $3j$, $6j$, CGC, F and R all have closed-form single-sum expressions — the Racah calculus [5], with the Wigner–Eckart framework and the $3j$ symmetry structure of [16] — and `racah` evaluates them in exact big-rational arithmetic carried as a signed square-rooted rational, rounding once at the end. The factorial-heavy sums are held in prime-factorised form to avoid big-integer blow-up, following `WignerSymbols.jl` [10]. The phase and normalisation conventions (Condon–Shortley for $3j$ and CGC; the `TensorKitSectors` [19] argument order for F ; $(-1)^{j_1+j_2-j_3}$ for R) are gauge and live in [docs/gauge.md §12](#).

9 Duality and the Frobenius–Schur indicator

The *dual* $\bar{a}$ of an irrep a carries the complex-conjugate representation $\overline{\pi^a}$; equivalently its highest weight is $-w_0\lambda$ with w_0 the longest Weyl group element. An irrep is *self-dual* if $\bar{a} \cong a$, which happens exactly when V_a carries a G -invariant non-degenerate bilinear form. By Fact 6.2 that form is unique up to scale, hence either symmetric or antisymmetric.

Fact 9.1 (Frobenius–Schur). For $a \in \text{Irr}(G)$ with character χ_a , define

$$\varkappa_a = \int_G \chi_a(g^2) dg \in \{+1, -1, 0\}. \quad (9)$$

Then $\varkappa_a = +1$ iff a is *real* (orthogonal: the invariant form is symmetric, and π^a is conjugate to a real representation); $\varkappa_a = -1$ iff a is *pseudoreal* (symplectic: the form is antisymmetric); and $\varkappa_a = 0$ iff a is *complex*, i.e. not self-dual. Bröcker–tom Dieck [14] §II.6; [7] §3.5.

In the categorical language of Section 12, $\varkappa_a$ is also the number relating the two cups/caps that trivialise $a \otimes \bar{a}$ (Kitaev [18] Appendix E); this is the reading under which it enters the F/R conventions.

For the families `racah` implements, (9) reduces to a combinatorial rule on the labels, so indicators are exact discrete data and never numerical results.

Fact 9.2 (Indicator of a self-dual irrep from the highest weight). For G connected compact simply connected and λ self-dual,

$$\varkappa_\lambda = (-1)^{\langle \lambda, 2\rho^\vee \rangle}, \quad 2\rho^\vee = \sum_{\alpha \in \Phi^+} \alpha^\vee \quad (10)$$

(the sum of the positive *coroots*). Bourbaki [20] Ch. VIII §7.

`racah` evaluates (10) in the ε -basis. The consequences it relies on: for $SU(2)$, $\varkappa = (-1)^{2j}$ (+1 for integer spin, -1 for half-integer); every tensor irrep of B_r and D_r is orthogonal (+1); and for C_r , $\varkappa = (-1)^{\sum_i \lambda_i}$. API: `su2_frobenius_schur`, `bcd::Irrep::frobenius_schur`, `{sun,bcd}::Irrep::dual` (for $SU(N)$ the dual reverses the Dynkin tuple, so self-duality is the palindrome condition and no separate indicator surface is exposed).

10 Tensor products, fusion multiplicities, outer multiplicity

By Fact 6.1 the tensor product of two irreps decomposes,

$$V_a \otimes V_b \cong \bigoplus_{c \in \text{Irr}(G)} \mathbb{C}^{N_{ab}^c} \otimes V_c, \quad N_{ab}^c = \dim \text{Hom}_G(V_c, V_a \otimes V_b) \in \mathbb{Z}_{\geq 0}, \quad (11)$$

where the second equality is Fact 6.2. N_{ab}^c is the *fusion multiplicity*. Taking dimensions in (11),

$$\sum_c N_{ab}^c d_c = d_a d_b. \quad (12)$$

For $SU(2)$ every $N_{ab}^c \in \{0, 1\}$ — the coupling is *multiplicity-free* — but for $SU(N \geq 3)$, $SO(N)$ and $Sp(2N)$ a given c can occur several times. When $N_{ab}^c > 1$ the triple (a, b, c) does *not* name a coupling channel; the pair (c, μ) with $\mu = 1, \dots, N_{ab}^c$ does. That index μ is the *outer multiplicity*, and it propagates into every downstream object: it is the trailing CGC axis and the four $[\mu, \nu, \kappa, \lambda]$ axes of an F block.

Fact 10.1 (Computing N_{ab}^c). N_{ab}^c is exact integer combinatorics.

- *A-series*: the Littlewood–Richardson rule expresses N_{ab}^c as a count of LR skew tableaux; equivalently it is the coefficient extracted from the product of Schur functions. [7] §A.1, §6.1.
- *All series*: the Brauer–Klimyk (Racah–Speiser) rule computes $\text{ch } V_a \cdot \text{ch } V_b$ by shifting the weights of V_a by $b + \rho$, folding each result back into the dominant chamber under the Weyl group, and accumulating the sign $\det(w)$ — weights fixed by a reflection cancel. [8] §24.4.

Implementations: **sun** uses the LR route via `SUNRepresentations.jl` [9]; **bcd** uses Brauer–Klimyk over Freudenthal weight multiplicities (8), with `QSpace` [3] and `GroupMath` [6] fixtures as cross-checks.


Convention 10.2 (Multiplicity ordering). Nothing in the mathematics orders the μ -copies: any basis of the N_{ab}^c -dimensional space $\text{Hom}_G(V_c, V_a \otimes V_b)$ will do. The order is therefore gauge. For $SU(N)$ it is the column order that the joint gauge-fixing of the whole nullspace block produces ([docs/gauge.md](#) §8); for $B/C/D$ it is discovery order in the sweep ([docs/gauge_soN.md](#) §9). Neither is an independent knob: “normalising” the multiplicity columns by sorting or reversing them is a specification deviation, not a refactor.

11 Clebsch–Gordan coefficients

Fix orthonormal bases $\{|a, m_a\rangle\}$ of each V_a and, for each channel $(a, b \rightarrow c, \mu)$, an isometric intertwiner $\iota_{ab}^{c\mu} : V_c \hookrightarrow V_a \otimes V_b$ realising (11). The *Clebsch–Gordan coefficients* (CGC) are its matrix elements: writing $|c, m_c; \mu\rangle := \iota_{ab}^{c\mu} |c, m_c\rangle$,

$$|c, m_c; \mu\rangle = \sum_{m_a, m_b} C_{a m_a; b m_b}^{c\mu} |a, m_a\rangle \otimes |b, m_b\rangle. \quad (13)$$

(The dependence of C on m_c is left implicit in the symbol, as in the code; the full index tuple is $C[m_a, m_b, m_c, \mu]$.)



The fusion vertex: the splitting intertwiner $\iota_{ab}^{c\mu}$ of (13), read bottom to top, whose matrix elements are the CGC.

11.1 Orthonormality and completeness

These are not conventions; they follow from Facts 6.2 and 6.3. Isometry of each $\iota_{ab}^{c\mu}$ and orthogonality of distinct channels give

$$\sum_{m_a, m_b} \overline{C_{a m_a; b m_b}^{c\mu}} C_{a m_a; b m_b}^{c'\mu'} = \delta_{cc'} \delta_{\mu\mu'} \delta_{m_c m_{c'}}, \quad (14)$$

and, since the channels together exhaust $V_a \otimes V_b$,

$$\sum_{c, \mu, m_c} C_{a m_a; b m_b}^{c\mu} \overline{C_{a m'_a; b m'_b}^{c\mu}} = \delta_{m_a m'_a} \delta_{m_b m'_b}, \quad (15)$$

whose trace is the dimension count (12). Equivalently: the matrix $M[(m_a, m_b), (m_c, \mu)] = C_{a m_a; b m_b}^{c\mu}$ is unitary once the column index runs over all (c, μ, m_c) .

racah runs (14) as a *generation gate*: a residual above tolerance is a typed error, never a silently returned coefficient.

11.2 Gauge freedom

The CGC are not determined by (a, b, c, μ) . Two choices are free:

- (a) the orthonormal basis of each V_x , i.e. a unitary U^x per irrep;
- (b) when $N_{ab}^c > 1$, how the μ -copies are oriented inside the isotypic component, i.e. a unitary $V_{ab}^c \in U(N_{ab}^c)$ mixing them.

Proposition 11.1 (The gauge group acting on CGC). *For unitaries U^a, U^b, U^c and $V_{ab}^c \in U(N_{ab}^c)$, the coefficients*

$$\tilde{C}_{a m_a; b m_b}^{c\mu} = \sum_{\mu', n_a, n_b, n_c} V_{ab, \mu\mu'}^c U_{m_a n_a}^a U_{m_b n_b}^b \overline{U_{m_c n_c}^c} C_{a n_a; b n_b}^{c\mu'} \quad (16)$$

satisfy (13)–(15) for the transformed bases and describe the same representation theory. This is the gauge freedom; (16) is the gauge transformation. Note that (b) is a genuinely per- (a, b, c) freedom — a unitary attached to each multiplicity slot — and is invisible whenever $N_{ab}^c \leq 1$.

Proposition 11.2 (What is gauge-invariant). *The orthogonal projector onto the c -isotypic component of $V_a \otimes V_b$,*

$$[P_{ab}^c]_{(m_a, m_b), (m'_a, m'_b)} = \left[\sum_{\mu} \iota_{ab}^{c\mu} (\iota_{ab}^{c\mu})^\dagger \right] = \sum_{\mu, m_c} C_{a m_a; b m_b}^{c\mu} \overline{C_{a m'_a; b m'_b}^{c\mu}}, \quad (17)$$

is unchanged by (16) up to conjugation by $U^a \otimes U^b$, and is completely invariant when the leg bases are held fixed. It is therefore the right object for cross-implementation comparison.

Freezing a gauge. “Fixing the gauge” means choosing, once and deterministically, a representative of each gauge orbit — as a specified function of exact combinatorial data (the ordered basis, a pivot rule, a sign rule, a descent order), never as “whatever the current code happens to output”. That distinction is the whole point of [docs/gauge.md](#) and [docs/gauge_soN.md](#): under a frozen specification, a value contradicting a stated rule is a *bug* by definition, and a consumer’s checkpointed coefficients stay valid across refactors. **racah** freezes:

- for $SU(N)$: the GT basis order, the coupling-pair column order, the nullspace tolerance and rank rule, the column-pivoted reduced-echelon pivot and tie-break, the positive-diagonal QR sign fix, and the reverse-lex descent order ([docs/gauge.md](#) §1–§8);

- for $B/C/D$: the Kronecker packing, the seed matrices, the QR gauge, the descending-weight sort and tie-break, the first-significant-entry block sign, the discovery-order multiplicity assignment, the canonical-parent rule, and the intertwiner alignment of rediscovered frames ([docs/gauge_soN.md](#) §1–§16);
- for $SU(2)$: the closed-form phase/normalisation conventions ([docs/gauge.md](#) §12).

Cross-checks against a *different* convention necessarily go through the gauge-invariant projector (17) or an explicit gauge-transformation harness — which is exactly how the QSpace CGC oracle for $B/C/D$ is built [3]. Where the reference convention is the same, the check is signed and element-wise: `SUNRepresentations.jl` [9] for $SU(N)$, and the `wigner-symbols` crate [11] for $SU(2)$.

12 Recoupling: $6j$, F -symbols, R -symbols

Fusing three irreps a, b, c into d can be bracketed two ways, through an intermediate e on the left or f on the right:

$$(V_a \otimes V_b) \otimes V_c \supset V_e \otimes V_c \supset V_d, \quad V_a \otimes (V_b \otimes V_c) \supset V_a \otimes V_f \supset V_d.$$

Both families of composites are bases of the same multiplicity space $\text{Hom}_G(V_d, V_a \otimes V_b \otimes V_c)$, so there is a unitary change of basis between them. That unitary is the *associator*, or F -symbol.

12.1 $SU(2)$: the $6j$ symbol

For $SU(2)$, where all multiplicities are 0 or 1, the change of basis is Racah’s $6j$ symbol $\left\{ \begin{smallmatrix} j_1 & j_2 & j_3 \\ j_4 & j_5 & j_6 \end{smallmatrix} \right\}$, which has a closed-form single-sum expression over products of factorials [5, 16]. It relates $(a \otimes b) \otimes c$ coupled through e to $a \otimes (b \otimes c)$ coupled through f , and `racah` evaluates it exactly, keyed by its Regge symmetry class. The F -symbol is then obtained from it by a fixed dimension-factor-and-phase dressing whose *argument order is itself gauge*; [docs/gauge.md](#) §12 is the authority for that formula. API: `wigner_3j`, `wigner_6j`, `su2_f_symbol`, `su2_r_symbol`.

12.2 General G : the F -symbol as a four-CGC contraction

For general G no closed form is available and F is *defined* by contracting four CGC over every magnetic index, so that only the multiplicity axes survive.

Convention 12.1 (Axis order and vertex assignment of an F block). A returned F block is a dense rank-4 array $[F_d^{abc}]_{e,f}[\mu, \nu, \kappa, \lambda]$, indexed by the two intermediate labels e, f and by four multiplicity indices, one per fusion vertex:

$$\underbrace{a \otimes b \rightarrow e}_{\mu \in [1, N_{ab}^e]}, \quad \underbrace{e \otimes c \rightarrow d}_{\nu \in [1, N_{ec}^d]}, \quad \underbrace{b \otimes c \rightarrow f}_{\kappa \in [1, N_{bc}^f]}, \quad \underbrace{a \otimes f \rightarrow d}_{\lambda \in [1, N_{af}^d]}.$$

So e (with μ, ν) belongs to the **left** bracketing $(a \otimes b) \otimes c$ and f (with κ, λ) to the **right** bracketing $a \otimes (b \otimes c)$; the axis lengths are $[N_{ab}^e, N_{ec}^d, N_{bc}^f, N_{af}^d]$ in that order, stored row-major, matching the `TensorKitSectors` [19] `GenericFusion` convention. The contraction wiring and this axis order are shared between the $SU(N)$ and $B/C/D$ engines by a single private core, so the convention cannot drift between families. [docs/gauge.md](#) is the normative authority.

With that assignment the definition `racah` implements is

$$[F_d^{abc}]_{e,f}[\mu, \nu, \kappa, \lambda] = \sum_{m_a, m_b, m_c, m_e, m_f} C_{a m_a; b m_b}^{e \mu} C_{e m_e; c m_c}^{d \nu} \Big|_{m_d} \overline{C_{b m_b; c m_c}^{f \kappa}} \overline{C_{a m_a; f m_f}^{d \lambda}} \Big|_{m_d}, \quad (18)$$

where the two CGC landing in d are evaluated at one *fixed* d -basis index m_d — `racah` takes the first, $m_d = 1$. Fixing rather than summing is not an approximation: by Fact 6.2 the result is independent of which d -state is used, which is itself a usable self-check. If any of the four vertex multiplicities vanishes the block is empty; the public wrappers report that as a typed *zero fusion channel* error rather than returning a zero block.

$$\begin{array}{c} a \quad b \quad c \\ \diagdown \quad \diagup \quad \diagup \\ \mu \quad e \quad \nu \\ \diagup \quad \diagdown \quad \diagdown \\ d \end{array} = \sum_{f, \kappa, \lambda} [F_d^{abc}]_{e, f} [\mu, \nu, \kappa, \lambda] \begin{array}{c} a \quad b \quad c \\ \diagdown \quad \diagup \quad \diagup \\ \kappa \quad f \quad \lambda \\ \diagup \quad \diagdown \quad \diagdown \\ d \end{array}$$

The F -move: the basis change computed by (18), with the four vertices labelled exactly as in Convention 12.1.

Remark 12.2 (Real CGC, elided conjugation). In `racah`'s frozen gauge the CGC of *every* family come out real-valued: the $SU(N)$ construction solves a real nullspace and gauge-fixes by a real echelon/QR pair, and the $B/C/D$ seeds are integer matrices carried through real orthogonalizations. The implementation therefore elides the conjugations in (18) and (20) — value-identically, not as an approximation. They are written here because they are part of the definition, and because a future complex gauge would need them back.

Unitarity of F in the combined indices holds by (14)–(15): arranging F as a matrix M with rows (e, μ, ν) and columns (f, κ, λ) ,

$$\sum_{f, \kappa, \lambda} [F_d^{abc}]_{e, f} [\mu, \nu, \kappa, \lambda] \overline{[F_d^{abc}]_{e', f} [\mu', \nu', \kappa, \lambda]} = \delta_{ee'} \delta_{\mu\mu'} \delta_{\nu\nu'}, \quad (19)$$

i.e. $MM^\dagger = \mathbb{I}$ ($MM^\top = \mathbb{I}$, M being real). This is shipped as a public self-check.

12.3 The R -symbol

The braiding exchanges the order of two fused irreps. Writing $\tau_{ab} : V_a \otimes V_b \rightarrow V_b \otimes V_a$ for the flip $v \otimes w \mapsto w \otimes v$, the composite $\tau_{ab} \circ \iota_{ab}^{c\mu}$ is an intertwiner $V_c \rightarrow V_b \otimes V_a$, hence a combination of the $\iota_{ba}^{c\nu}$:

$$\tau_{ab} \circ \iota_{ab}^{c\mu} = \sum_{\nu} [R_c^{ab}]_{\mu\nu} \iota_{ba}^{c\nu}, \quad \text{i.e.} \quad [R_c^{ab}]_{\mu\nu} = \sum_{m_a, m_b} C_{a m_a; b m_b}^{c\mu} \Big|_{m_c} \overline{C_{b m_b; a m_a}^{c\nu} \Big|_{m_c}}, \quad (20)$$

again at a single fixed m_c (the first), for the same Schur-lemma reason as in (18). Note the index order in the second factor: its *first* magnetic slot is m_b , because it is a CGC of the coupling $b \otimes a \rightarrow c$.

$$\begin{array}{c} b \quad a \\ \diagdown \quad \diagup \\ \mu \\ \diagup \quad \diagdown \\ c \end{array} = \sum_{\nu} [R_c^{ab}]_{\mu\nu} \begin{array}{c} b \quad a \\ \diagdown \quad \diagup \\ \nu \\ \diagup \quad \diagdown \\ c \end{array}$$

The R -move: the braiding $\tau_{ab} \circ \iota_{ab}^{c\mu} = \sum_{\nu} [R_c^{ab}]_{\mu\nu} \iota_{ba}^{c\nu}$ of (20).

R_c^{ab} is an $N_{ab}^c \times N_{ba}^c$ unitary — *not* square in general, though $N_{ab}^c = N_{ba}^c$ always holds — with rows indexed by μ and columns by ν . In the multiplicity-free case it is a single phase, and for $SU(2)$ it is $(-1)^{j_a + j_b - j_c}$ on an admissible triangle and exactly 0 otherwise (the zero mirrors $N_{ab}^c = 0$ and stops a caller multiplying a spurious sign into a forbidden channel). API: `sun::{\f_symbol, r_symbol}, bcd::{\f_symbol, r_symbol}`.

Remark 12.3 (No $6j$ outside $SU(2)$). The recoupling surface for the generated families is F and R only. There is no $6j$ function for $SU(N)$ or $B/C/D$: with outer multiplicity present the $6j$'s scalar shape is not the natural object, and F carries the same information with the axes made explicit.

13 The pentagon and hexagon identities

F and R are not free data. They are the associator and braiding of a unitary braided fusion category, and as such must satisfy the coherence axioms of Mac Lane's theorem: the *pentagon* (associativity of four-fold fusion) and the two *hexagons* (compatibility of braiding with fusion). The axioms in the form used here are Etingof–Gelaki–Nikshych–Ostrik [17] §§2.1, 8.1; the index-carrying physicist's form, with the same conventions this crate follows, is Kitaev [18] Appendix E.

Both are stated below in the *exact index convention of the crate*: every F carries its four multiplicity axes in the order of Convention 12.1, and $[F_w^{xyz}]_{u,v}[\cdot, \cdot, \cdot, \cdot]$ means the block for fusing x, y, z into w through u (left) or v (right).

13.1 Pentagon

Let a, b, c, d be irreps and quantify over every

$$f \in a \otimes b, \quad h \in c \otimes d, \quad g \in f \otimes c, \quad i \in b \otimes h, \quad e \in (g \otimes d) \cap (a \otimes i).$$


Then, for every value of the six *free* multiplicity indices

$$\lambda \in [1, N_{fc}^g], \quad \mu \in [1, N_{gd}^e], \quad \nu \in [1, N_{cd}^h], \quad \kappa \in [1, N_{ab}^f], \quad \rho \in [1, N_{bh}^i], \quad \sigma \in [1, N_{ai}^e],$$

the pentagon identity is

$$\begin{aligned} & \sum_{\tau} [F_e^{fcd}]_{g,h}[\lambda, \mu, \nu, \tau] [F_e^{abh}]_{f,i}[\kappa, \tau, \rho, \sigma] \\ &= \sum_{j \in b \otimes c} \sum_{\alpha, \beta, \tau'} [F_g^{abc}]_{f,j}[\kappa, \lambda, \alpha, \beta] [F_e^{ajd}]_{g,i}[\beta, \mu, \tau', \sigma] [F_i^{bcd}]_{j,h}[\alpha, \tau', \nu, \rho], \end{aligned} \quad (21)$$

where the contracted multiplicity indices run over $\tau \in [1, N_{fh}^e]$ on the left and, on the right, $\alpha \in [1, N_{bc}^j]$, $\beta \in [1, N_{aj}^g]$, $\tau' \in [1, N_{jd}^i]$. A label tuple for which any free-index range is empty is vacuous and skipped. This is a condition on F alone — associativity of four-fold fusion.



13.2 Hexagons

Quantify over every $e \in c \otimes a$, $f \in c \otimes b$, $d \in (e \otimes b) \cap (a \otimes f)$, and every value of the free indices

$$\alpha \in [1, N_{ca}^e], \quad \beta \in [1, N_{eb}^d], \quad \mu \in [1, N_{bc}^f], \quad \nu \in [1, N_{af}^d].$$

The first hexagon is

$$\begin{aligned} \sum_{\lambda, \gamma} [R_e^{ca}]_{\alpha\lambda} [F_d^{acb}]_{e,f}[\lambda, \beta, \gamma, \nu] [R_f^{cb}]_{\gamma\mu} \\ = \sum_{g \in a \otimes b} \sum_{\delta, \sigma, \psi} [F_d^{cab}]_{e,g}[\alpha, \beta, \delta, \sigma] [R_d^{cg}]_{\sigma\psi} [F_d^{abc}]_{g,f}[\delta, \psi, \mu, \nu], \end{aligned} \quad (22)$$

with contracted indices $\lambda \in [1, N_{ac}^e]$, $\gamma \in [1, N_{cb}^f]$ on the left and $\delta \in [1, N_{ab}^g]$, $\sigma \in [1, N_{cg}^d]$, $\psi \in [1, N_{gc}^d]$ on the right. The second hexagon is (22) with the three braidings replaced by the opposite ones and conjugated,

$$R_e^{ca} \rightsquigarrow \overline{R_e^{ac}}, \quad R_f^{cb} \rightsquigarrow \overline{R_f^{bc}}, \quad R_d^{cg} \rightsquigarrow \overline{R_d^{gc}},$$

which is the $R \mapsto R^{-1}$ form of the axiom written out; since R is real in this gauge (Remark 12.2) the conjugations are elided in the implementation and the two hexagons differ only in the argument order of the three R blocks.

Remark 13.1 (Status in the implementation). (21) and (22) are shipped as *public self-checks* and additionally run as *generation gates*: a violation beyond tolerance is a typed error, never a silently returned coefficient. They are the strongest available consistency statement about numerically constructed CGC, because they constrain the whole coefficient set jointly rather than one block at a time. The equations, index conventions and the generic-multiplicity branch follow TensorKitSectors [19]; see [docs/references.md](#) for the `file:symbol`-level rows.

Remark 13.2 (Gauge covariance of the axioms). A gauge transformation (16) changes F and R but preserves (21) and (22): the transformation acts on F by conjugation with the V 's on each of the four multiplicity slots and the U 's cancel. So passing the pentagon and hexagon certifies *consistency*, not agreement with any particular convention. Convention agreement is a separate claim, and it is settled by the oracles of Section 14 against the pinned reference implementations.

14 What is exact, and what is verification-gated

The note's objects split cleanly into combinatorial data and solved data, and the implementation's contract follows that split.

Exact / combinatorial. Irrep labels and dominance; the Weyl dimension (7); duals; Frobenius–Schur indicators (9); fusion multiplicities (11); weight systems (8); GT pattern enumeration and basis ordering; and, for $SU(2)$ only, the full $3j/6j/CGC/F/R$ values, carried as signed square-rooted rationals with a single final rounding.

Numeric / verification-gated. For the generated families the CGC — and the F , R contracted from them by (18), (20) — come from a nullspace solve and are floating point. Their *values* are finite precision, but the *gauge* fixing them is a deterministic function of exact combinatorial data, and every generation runs the orthonormality (14), unitarity (19), pentagon (21) and hexagon (22) checks before returning.

So “exact” in this project is a claim about **structure, gauge determinism, and verification**, not about symbolic algebraic-number coefficient values.

Independent oracles. Consistency checks cannot detect a systematic convention divergence (see the gauge-covariance remark in Section 13), so each family is additionally checked against an implementation that is not this one:

Family	Oracle	Nature of the comparison
$SU(2)$	<code>wigner-symbols</code> crate v0.5.1 [11]	exhaustive over its label domain
$SU(N)$	<code>SUNRepresentations.jl</code> v0.4.0 [9]	signed, element-wise, including the μ order
$SU(N)$	<code>GroupMath</code> v1.1.3 [6]	products and multiplicities
$SO(N)$, $Sp(2N)$	<code>QSpace</code> v4 [3]	projector battery (17)
$B/C/D$ dims	<code>QSpace</code> v4 [3]	exact dimension agreement

The $SU(2)$ and $SU(N)$ comparisons are signed and element-wise because `racah` adopts those references’ conventions outright; the $B/C/D$ comparison goes through the gauge-invariant projector (17) because `racah`’s $B/C/D$ gauge deliberately deviates from `QSpace`’s in documented, enumerated ways (`docs/gauge_soN.md`).

References

- [1] A. Alex, M. Kalus, A. Huckleberry, J. von Delft, “A numerical algorithm for the explicit calculation of $SU(N)$ and $SL(N, \mathbb{C})$ Clebsch–Gordan coefficients,” *J. Math. Phys.* **52**, 023507 (2011). [doi:10.1063/1.3521562](https://doi.org/10.1063/1.3521562); [arXiv:1009.0437](https://arxiv.org/abs/1009.0437).
- [2] A. Weichselbaum, “Non-abelian symmetries in tensor networks: A quantum symmetry space approach,” *Ann. Phys.* **327**, 2972–3047 (2012). [doi:10.1016/j.aop.2012.07.009](https://doi.org/10.1016/j.aop.2012.07.009); [arXiv:1202.5664](https://arxiv.org/abs/1202.5664).
- [3] A. Weichselbaum, “QSpace — An open-source tensor library for Abelian and non-Abelian symmetries,” *SciPost Phys. Codebases* **40** (2024). [doi:10.21468/SciPostPhysCodeb.40](https://doi.org/10.21468/SciPostPhysCodeb.40); [arXiv:2405.06632](https://arxiv.org/abs/2405.06632). `racah` cites `QSpace` v4 source at revision `dd2cc7e`.
- [4] I. M. Gelfand, M. L. Tsetlin, “Finite-dimensional representations of the group of unimodular matrices,” *Dokl. Akad. Nauk SSSR* **71**, 825–828 (1950) (Russian). No DOI (predates DOI registration).
- [5] G. Racah, “Theory of Complex Spectra. II,” *Phys. Rev.* **62**, 438–462 (1942). [doi:10.1103/PhysRev.62.438](https://doi.org/10.1103/PhysRev.62.438).
- [6] R. M. Fonseca, “GroupMath: A Mathematica package for group theory calculations,” *Comput. Phys. Commun.* **267**, 108085 (2021). [doi:10.1016/j.cpc.2021.108085](https://doi.org/10.1016/j.cpc.2021.108085); [arXiv:2011.01764](https://arxiv.org/abs/2011.01764). `racah` fixtures use `GroupMath` v1.1.3.
- [7] W. Fulton, J. Harris, *Representation Theory: A First Course*, Graduate Texts in Mathematics **129**, Springer (1991). ISBN 0-387-97495-4. [doi:10.1007/978-1-4612-0979-9](https://doi.org/10.1007/978-1-4612-0979-9).
- [8] J. E. Humphreys, *Introduction to Lie Algebras and Representation Theory*, Graduate Texts in Mathematics **9**, Springer (1972). ISBN 978-0-387-90053-7. [doi:10.1007/978-1-4612-6398-2](https://doi.org/10.1007/978-1-4612-6398-2).
- [9] `SUNRepresentations.jl` (v0.4.0), a Julia implementation of [1]. <https://github.com/QuantumKitHub/SUNRepresentations.jl>.
- [10] `WignerSymbols.jl` (v2.0.0), J. Haegeman. <https://github.com/Jutho/WignerSymbols.jl>.
- [11] `wigner-symbols` Rust crate (v0.5.1), P. Ruffwind. <https://crates.io/crates/wigner-symbols>.

- [12] A. I. Molev, “Gelfand–Tsetlin bases for classical Lie algebras,” in *Handbook of Algebra* **4**, M. Hazewinkel (ed.), Elsevier (2006), pp. 109–170. [doi:10.1016/S1570-7954\(06\)80006-9](https://doi.org/10.1016/S1570-7954(06)80006-9); [arXiv:math/0211289](https://arxiv.org/abs/math/0211289).
- [13] R. Slansky, “Group theory for unified model building,” *Phys. Rep.* **79**, 1–128 (1981). [doi:10.1016/0370-1573\(81\)90092-2](https://doi.org/10.1016/0370-1573(81)90092-2). Congruency classes: §5, p. 37; table notes for Table 36 ($SO(8)$) and Table 41 ($SO(10)$).
- [14] T. Bröcker, T. tom Dieck, *Representations of Compact Lie Groups*, Graduate Texts in Mathematics **98**, Springer (1985). ISBN 978-3-540-13678-1. [doi:10.1007/978-3-662-12918-0](https://doi.org/10.1007/978-3-662-12918-0). *Added for this note:* Haar averaging and complete reducibility (§II.1–II.4), Peter–Weyl and Schur orthogonality (§III.3), and the Frobenius–Schur indicator (§II.6).
- [15] G. E. Baird, L. C. Biedenharn, “On the Representations of the Semisimple Lie Groups. II,” *J. Math. Phys.* **4**, 1449–1466 (1963). [doi:10.1063/1.1703926](https://doi.org/10.1063/1.1703926). *Added for this note:* the explicit closed-form matrix elements of the $\mathfrak{gl}_N$ generators in the Gelfand–Tsetlin basis, which Fact 8.2 states and the $SU(N)$ construction consumes.
- [16] E. P. Wigner, *Group Theory and its Application to the Quantum Mechanics of Atomic Spectra*, Academic Press (1959). *Added for this note:* the Wigner–Eckart framework and the $3j$ symmetry structure underlying the $SU(2)$ closed forms.
- [17] P. Etingof, S. Gelaki, D. Nikshych, V. Ostrik, *Tensor Categories*, Mathematical Surveys and Monographs **205**, American Mathematical Society (2015). ISBN 978-1-4704-2024-6. [doi:10.1090/surv/205](https://doi.org/10.1090/surv/205). *Added for this note:* the fusion-category axioms — pentagon (§2.1) and hexagon (§8.1) — as the source for (21) and (22).
- [18] A. Kitaev, “Anyons in an exactly solved model and beyond,” *Ann. Phys.* **321**, 2–111 (2006). [doi:10.1016/j.aop.2005.10.005](https://doi.org/10.1016/j.aop.2005.10.005); [arXiv:cond-mat/0506438](https://arxiv.org/abs/cond-mat/0506438). *Added for this note:* Appendix E, the index-carrying (F , R , Frobenius–Schur) conventions in the form this crate’s conventions follow.
- [19] TensorKitSectors.jl, the categorical-symmetry layer of TensorKit. <https://github.com/QuantumKitHub/TensorKitSectors.jl>. *Added for this note as a numbered entry* ([docs/references.md](#) cites it inline at `file:symbol` level): the pinned source for the F/R contraction wiring and axis order (`sectors.jl:Fsymbol_from_fusientensor`), the pentagon/hexagon index conventions (`sectors.jl:pentagon_equation`, `hexagon_equation`), and the $SU(2)$ F/R /Frobenius–Schur conventions (`irreps/su2irrep.jl`). It is the **pinned convention oracle** for $SU(2)$: where the mathematics admits several conventions, `racah`’s choice is by definition TensorKitSectors’ choice.
- [20] N. Bourbaki, *Lie Groups and Lie Algebras, Chapters 4–6*, Elements of Mathematics, Springer (2002). ISBN 978-3-540-42650-9. [doi:10.1007/978-3-540-89394-3](https://doi.org/10.1007/978-3-540-89394-3). *Added for this note:* the simple-root numbering (Plates I–IX) that fixes the order of the Dynkin label tuple in the API; Ch. VI §4 (Theorem 3 and the plates) for the Killing–Cartan classification, the nine Dynkin diagrams, the low-rank identifications (§4.4), and the tabulated P/Q of each family.
- [21] A. W. Knap, *Lie Groups Beyond an Introduction*, 2nd ed., Progress in Mathematics **140**, Birkhäuser (2002). ISBN 978-0-8176-4259-4. [doi:10.1007/978-1-4757-2453-0](https://doi.org/10.1007/978-1-4757-2453-0). *Added for this note:* existence of a compact real form of a complex semisimple Lie algebra (Theorem 6.11) and its uniqueness up to conjugacy (Corollary 6.20) — the step from the classification of algebras to the compact groups `racah` computes for.