

Notes on Rubik's abstract polytopes

Utensil Song

October 9, 2026

These are the accompanying mathematical notes for the experimental research project [FRAP](#), *Formalizing Rubik's Abstract Polytopes*. The project asks how the general mathematics of Rubik puzzles can explain the concrete practice of solving them. Mathematical exposition, formalization in Lean, and executable algorithms are intended to develop together: understanding a puzzle should help us solve it, while trying to solve it should sharpen the mathematics.

The motivation grew out of decades of playing the $3 \times 3 \times 3$ cube, followed by the $2 \times 2 \times 2$ and $4 \times 4 \times 4$. Learning formulas raised questions about why they work, how different methods organize a solve, and how a computer might find a better solution. Those questions led back through group theory and the general $N \times N \times N$ family to Marchetti's *Rubik's Abstract Polytopes* [[Marchetti\(2025\)](#)]. These notes deliberately reverse that journey: begin with general mathematics, develop increasingly concrete models, and return to solution methods and play with a deeper understanding of the familiar puzzle.

A first concrete question guides the route: when a turn acts on one facet, which positions move, and what incident labels travel with them? Starting with abstract polytopes lets us answer that question in two stages. We first record incidence and shape symmetry without choosing lengths or angles, then apply a local facet action and inspect its effect on a tetrahedron and a cube. This separation keeps the shape data distinct from the move data while giving each abstraction a visible puzzle consequence. Passing to layered cubes then asks how the pieces, their orientations and the allowed turns change with N ; the $3 \times 3 \times 3$ cube provides a setting in which to work out the consequences in detail. Group actions connect these descriptions. They help us ask what a move changes, what a sequence preserves, which configurations are reachable, and how a solution can be assembled from smaller tasks. The additional structure in each model must be explained when the models are compared.

The return to solving has two closely related purposes. For a person, a useful formula connects a recognizable pattern with a manageable sequence of moves.

Understanding its effect and what it preserves should help explain methods and the choices within a solve; recognition, memory and ease of execution matter alongside move count. For a computer, the question is how to find solutions under explicit objectives, including shortest solutions for a chosen set of moves and costs. Producing a solution, proving that a method solves every state in its domain, and establishing optimality are different achievements. Scrambling, replay and interactive exploration bring these questions back to actual use.

The puzzles also lead into mathematics beyond solving. Local orientation data raise questions about global constraints, wreath products and transfer. Transporting labels suggests comparisons with coloring obstructions and branched coverings [Izmestiev and Joswig(2002)]. Random scrambles and functions used to guide search lead to probability and representation theory [Pan and Kondor(2021)]. These are directions for inquiry, with their own constructions and hypotheses; the aim is to understand the connections through examples and arguments as well as the puzzles themselves.

Scholarly sources, explicit calculations and formalization are meant to inform one another. Lean can express mathematical models and implement computations whose stated properties can be proved. Missing hypotheses or awkward conventions exposed by that work should feed back into the explanations here. FRAP is at an early research stage, and the companion repository does not yet contain a formalized puzzle model or solver.

1 How the notes are organized

The developed mathematics follows one practical question: what data must a move carry so that we can later compose moves and test reachability? **Incidence, symmetry and moves** answers it in order. First, **polytopes, flags and facet symmetry** describe the shape and show their role in triangle and tetrahedron examples. Next, **facet moves** turn a local facet rotation into an action on incidence positions, with tetrahedron and cube calculations that make the construction concrete. Only then do **positions, orientations and global constraints** separate a piece's destination from transport in its local frame, while **transfer and the effect of dimension** connect that transport with a group homomorphism and compute the relevant local abelianizations.

The cube examples above connect the general mathematics with solving and play. The later accounts use those concrete states and moves as their starting point:

The transition between the abstract facet action and the fixed-frame cube coordinates is recorded in [how the cube models connect](#). Read it before the all- N law

if the distinction between incidence positions, cubie fibres and observations is new.


1. **General Rubik puzzles:** the incidence actions, framed orientations and transfer developed below.
2. **The [layered-cube family](#):** piece types, layer turns and conventions as N varies.
3. **The [layered-cube law](#):** all- N state coordinates and a constructive characterization of reachability.
4. **The [complete three-layer cube](#):** explicit outer generators, impossible states, the standard count and a worked cleanup.
5. **[Human methods and computer search](#):** recognizable cases, algebraic effects, usability, move costs and optimality. The [complete layer-first method](#) supplies one full human construction, and [computer solving and efficiency](#) supplies exact small search, admissible bounds and subgroup reduction.
6. **[From move sequences to play](#):** notation, scrambling, recognition, replay and complete solves, including the [worked replay and projected lower bound](#).

[Further connections](#) relate the developed orientation and transfer calculations to [label transport and topology](#) and [scrambling and search functions](#). The [quotient-compatible move account](#) gives the criterion for an identification of observations to inherit a puzzle action. Together these results compare transport, randomization and information loss around the same concrete move.

1.1 Five encounters with a cube


Put a solved three-layer cube on a table and turn its right face once. The move is familiar, but it raises concrete puzzles: which labelled pieces moved, why can a visually plausible arrangement be impossible, and how can another person undo the turn without watching the first hand? Begin with these experiences, then introduce only the distinctions needed to explain them.

1.1.1 One turn is larger than the visible face



The shaded strip is only the visible part of the layer: the same stickers cross the front, right, back, top, and bottom faces; viewed from above, the top-face part is the corresponding rightmost strip. A whole-cube rotation would change the viewing frame, whereas this face move keeps the six centres fixed and changes one layer. The [layered-cube family](#) makes this distinction for every N ; the [complete three-layer cube](#) spells it out for the six outer generators.

1.1.2 A sticker needs a home and a frame




A player sees colours; a proof must also track each sticker's cubie and transported local frame. [How the cube models connect](#) follows one sticker through incidence position, cubie fibre, and fixed-frame coordinate. A colour quotient forgets labels, while a larger cube adds positions, so the state space changes with the question.

1.1.3 A plausible state can still be impossible


Take a solved cube and twist one corner in place. Every sticker still has a cubie-sized home, and the three colours at that corner look locally consistent. But

the corner orientations now have sum 1 modulo 3, whereas every face turn preserves sum 0. No sequence of legal turns can produce this state.




The [complete three-layer cube](#) derives the matching edge and permutation restrictions from the six generators. The calculation is useful at the table: it tells a solver when a scrambled-looking object was assembled illegally, and it tells a proof which state space is worth searching.

A human algorithm has to preserve these legal states while it borrows pieces from completed areas.



A formula often disturbs pieces that were already placed. The disturbance is acceptable only when the setup is undone and the completed pieces return. A working slot, a trigger, and an inverse setup make that promise visible. The [complete layer-first method](#) repeats this pattern through the cross, the first layer, the middle layer, and the last layer. Its completeness depends on the preservation and termination checks, not on the formula's familiarity.

1.1.4 Undoing a scramble checks the notation




The [worked replay and projected lower bound](#) fixes the viewing convention, records the ten tokens, and applies their inverse. Sticker replay returns to solved. A corner-only projected search finds no solution of depth at most two, so the evidence is the bounded statement

$$3 \leq d_{\text{HTM}} \leq 10. \quad (1.1.4.1)$$

The inverse word is an upper-bound witness; the projected search is a lower-bound witness. The [computer account](#) explains why these checks do not prove that ten is shortest. A state, a word, and a replay log are enough to test the notation before asking for a general solver.

1.1.5 Changing one ingredient changes the problem

Once the move, its frame, its local repair, and its replay are explicit, three variations separate the effects of enlarging the cube, changing the shape, and forgetting labels.



The [layered-cube family](#) develops the first variation; the incidence and facet construction in [incidence, symmetry and moves](#) develops the second; and [which moves survive a quotient?](#) gives the descent condition for the third. [Framed orientations](#), [label transport and topology](#), and [scrambling, spectra and search](#) then study the extra information each variation introduces or forgets. Their hypotheses stay attached to the concrete move that prompted the question.

A face turn can therefore be drawn, named, replayed, and checked before it is generalized. The resulting state maps and invariants support later formal certification or measured comparisons, while the examples here make clear which claim each later programme would have to establish.

2 From incidence to moves

To describe a puzzle move without committing to coordinates, we need an incidence record that says which pieces and labels belong together. Faces and incidence define an abstract polytope without choosing lengths or an embedding [Schulte(2017), Section 2]. Regularity supplies shape symmetries; the Rubik construction applies facet rotations locally to incidence positions [Marchetti(2025), Definitions 5–6]. The sections below first make the shape visible in triangle and tetrahedron examples, then build this action and calculate its tetrahedron and cube consequences.

2.1 Faces, flags, and the symmetry of a shape

A local turn needs a shape-level answer to two questions: which faces are incident, and which symmetries preserve those incidences? An abstract polytope retains the incidence of vertices, edges and faces without prescribing distances or angles. Schulte’s flag formulation [Schulte(2017), Section 2] leads from these incidences to symmetries and to the extension of a facet symmetry used in **local puzzle moves**.

Definition 2.1.1 (Abstract polytopes and flags [Schulte(2017), Section 2])

An **abstract n -polytope**, for $n \geq 0$, is a **partially ordered set** P of **faces** with a least face F_{-1} and greatest face F_n . Every chain extends to a maximal chain, called a **flag**, with exactly $n + 2$ faces. Faces in each flag have **ranks** $-1, 0, \dots, n$, consistently across flags. Faces other than the least and greatest are *proper*; vertices, edges and facets have ranks $0, 1, n - 1$.

The **diamond condition** requires exactly two faces of rank i between faces $F < G$ of ranks $i - 1$ and $i + 1$, for $0 \leq i < n$. Two flags are *adjacent* if they differ in exactly one proper face. **Strong flag connectivity** requires any two flags to be joined by adjacent flags, each containing their intersection.

Thus, for each $0 \leq i < n$, each flag has a unique adjacent flag differing in rank i . Strong flag connectivity restricts paths to these adjacent flags, each containing the endpoint flags’ intersection; arbitrary paths through the least face would connect every pair without imposing an axiom. These conditions prescribe incidence without choosing lengths or angles.

A rank-zero polytope is a two-element chain. A rank-one polytope has two vertices between its least and greatest faces. These low-rank cases occur as **sections** of higher-rank polytopes.

Definition 2.1.2 (Sections and vertex figures


[Schulte(2017), Section 2])

Let P be a **abstract polytope**. For **faces** $F \leq G$ of **ranks** i, j , its **section** G/F is the interval $\{H : F \leq H \leq G\}$, with ranks shifted so that F has rank -1 . For $F < G$, it is an **abstract polytope** of rank $j - i - 1 \geq 0$. For $F = G$, the singleton section is called the rank -1 polytope. These low ranks allow all intervals, including vertex figures of polygons and intervals between consecutive faces.

The section F_n/v at a vertex v is its **vertex figure**. Its vertices are the edges through v ; its faces record the incidences above v . A facet H is viewed as the section H/F_{-1} .

Example 2.1.3 (Flags of a triangle and a tetrahedron)

The face-poset construction of [Schulte(2017), Section 2] gives the following elementary examples of **flags**. A triangle has the empty face, three vertices, three edges and the whole triangle, of **ranks** $-1, 0, 1, 2$. The **flag** $\emptyset < 1 < 12 < 123$ has rank-zero neighbor $\emptyset < 2 < 12 < 123$ and rank-one neighbor $\emptyset < 1 < 13 < 123$. There are $3 \cdot 2 = 6$ **flags**.



Each vertex in this flag graph records the vertex and edge of one flag; the common empty and greatest faces are suppressed. Edge labels give the rank changed by adjacency.

A tetrahedron has all subsets of $\{1, 2, 3, 4\}$ as **faces**, ordered by inclusion; a nonempty face with k vertices has rank $k - 1$. Its flag $\emptyset < 1 < 12 < 123 < 1234$ exhibits all five **ranks**. The **vertex figure** at 1 is a triangle: its vertices correspond to the three incident edges. A cube also has a triangular vertex figure, although its facets are squares. Local and global shapes carry different information.

The triangle and tetrahedron calculation now supplies a concrete check: the flag graph records which flags are adjacent; a vertex figure records the local neighborhood seen from one vertex. With those objects in hand, the next steps ask which incidence-preserving maps act on all flags and how a symmetry of one facet can be extended.

Definition 2.1.4 (Automorphisms and regularity

[Schulte(2017), Section 2])

An **automorphism** of an **abstract polytope** P is a bijection $a : P \rightarrow P$ satisfying $F \leq G \iff a(F) \leq a(G)$. Automorphisms form a group $\Gamma(P)$ under composition; in ab , the map b acts first. They preserve **ranks** and act on **flags** face by face.

The polytope is **regular** if this action is transitive on **flags**. The tetrahedron is regular: a permutation of its four vertices sends any flag to any other.

The underlying order automorphism is represented by `OrderIso` with equal source and target. This declaration does not impose the polytope axioms or assert regularity.

Lemma 2.1.5 (A flag determines an automorphism

[Schulte(2017), Section 2, p. 4])

An **automorphism** of an **abstract polytope** fixing every face of one **flag** is the identity. Consequently the automorphism group of a **regular** polytope acts freely and transitively on its **flags**.

The cited paragraph preceding Lemma 2.1 states simple transitivity for **regular** polytopes. The elementary argument below also proves freeness without assuming regularity.

Proof. For a **flag** adjacent to a pointwise fixed flag, only one rank has changed. The unchanged neighboring ranks bound a **diamond condition** with two possible faces. One is fixed, so the other is fixed too. The adjacent flag is therefore fixed pointwise. Propagating along the connected flag graph fixes every flag. Every face belongs to a flag, so every face is fixed. Transitivity is the definition of being **regular**. $\square$

Remark 2.1.6 (Distinguished generators and the intersection property)

Choose a base **flag** Φ in a **regular abstract polytope** P of rank n . By **flag freeness**, there is a unique **automorphism** ρ_i taking Φ to its rank- i neighbor, for $0 \leq i < n$. The **diamond condition** and uniqueness give $\rho_i^2 = 1$. For $|i - j| > 1$, the two rank changes are independent; $\rho_i \rho_j$ and $\rho_j \rho_i$ agree on Φ , hence are equal. Connectivity shows that the ρ_i reach every **flag**, so they generate $\Gamma(P)$. These are the **distinguished generators** of [Schulte(2017), Section 5, pp. 13–14].

For $I \subseteq \{0, \dots, n-1\}$, put $G_I = \langle \rho_i : i \in I \rangle$. They also satisfy

$$G_I \cap G_J = G_{I \cap J}. \quad (2.1.1)$$

This **intersection property** is part of the **string C-group** description [Schulte(2017), Section 5, equation (18)]. Involution and nonadjacent commutation relations alone do not reconstruct a polytope; the group-theoretic characterization includes this additional condition [Schulte(2017), Section 5, p. 14].

Lemma 2.1.7 (Extending a symmetry of one facet)

Let H be a facet of a **regular abstract polytope** P of rank $n \geq 2$. Restriction gives an isomorphism

$$\text{Stab}_{\Gamma(P)}(H) \cong \Gamma(H/F_{-1}). \quad (2.1.2)$$

In particular the facet is **regular**, and each of its **automorphisms** has a unique extension fixing H . This is a consequence of the flag action in [Schulte(2017), Section 2 and Lemma 2.1]; the argument follows.

Proof. Extend a flag of the **facet section** by the greatest face of P . A stabilizer element restricting to the identity fixes this whole flag and is the identity by **flag freeness**. Restriction is therefore injective.

Given a facet **automorphism** α , the **regular** action supplies an **automorphism** of P taking the extended flag to the extension of its image under α . It fixes H . Its restriction agrees with α on the chosen facet flag, so flag freeness in the facet makes these restrictions equal. This proves surjectivity. The same use of same **regular** action with any two facet flags shows that the facet is **regular**. $\square$

2.2 Turning a facet into a puzzle move

Once the shape symmetries are known, the next question is local: how can a turn change only the positions incident to one chosen facet while carrying their labels along? A facet rotation extends to a symmetry of the whole shape, but a puzzle move applies it only at positions incident to the chosen facet. Marchetti's construction retains both entries of each incidence pair [Marchetti(2025), Definitions 5–6]. We use finite regular abstract polytopes of rank at least three; the **extension lemma** itself does not require finiteness.

Definition 2.2.1 (The rotation subgroup [Marchetti(2025), Definition 4])

Let Q be a **regular abstract polytope** of rank $m \geq 2$. Choose a base **flag** and its **distinguished generators** (involutions) $\rho_0, \dots, \rho_{m-1}$. Its **rotation subgroup** is

$$\Gamma^+(Q) = \langle \rho_{j-1}\rho_j : 1 \leq j < m \rangle. \quad (2.2.1)$$

The index starts at 1: for a polygon, the single adjacent pair $\rho_0\rho_1$ generates its rotations. For triangles and squares these are the usual cyclic rotation groups.

Equivalently, $\Gamma^+(Q)$ consists of the elements represented by even-length words in the distinguished involutions. Indeed every pair $\rho_i \rho_j$ is a product of adjacent pairs or their inverses, by cancellation along the index interval. Conjugating an even word by a distinguished involution again gives an even word, so this subgroup is normal. A change of base flag conjugates the distinguished generators and therefore leaves the subgroup unchanged. We do not require it to have index two for every abstract polytope.

Definition 2.2.2 (Incidence positions and facet moves
[Marchetti(2025), Definitions 5–6])

Let P be a finite **regular abstract polytope** of rank $n \geq 3$, with least and greatest faces F_{-1}, F_n . Write $P^\circ = P \setminus \{F_{-1}, F_n\}$. The **incidence positions** are

$$\mathcal{R}(P) = \{(F, G) \in P^\circ \times P^\circ : F \leq G\}. \quad (2.2.2)$$

Choose a facet H and an element α of the **rotation subgroup** $\Gamma^+(H/F_{-1})$. The **facet extension lemma** gives its unique extension $\hat{\alpha} \in \Gamma(P)$ fixing H . The corresponding **facet move** is

$$\mu_{H,\alpha}(F, G) = \begin{cases} (\hat{\alpha}(F), \hat{\alpha}(G)), & F \leq H, \\ (F, G), & F \not\leq H. \end{cases} \quad (2.2.3)$$

The first entry decides where the shape automorphism is applied; the second entry is transported with it. The facet is part of the move's name: one shape **automorphism** may stabilize several facets and give different local moves. Each move is a **permutation** of $\mathcal{R}(P)$.

Lemma 2.2.3 (The inverse of a facet move)

Let P be a finite **regular abstract polytope** of rank $n \geq 3$, H a facet and $\alpha \in \Gamma^+(H/F_{-1})$ a facet rotation. The **facet move** $\mu_{H,\alpha}$ is a **permutation** of $\mathcal{R}(P)$, with

$$\mu_{H,\alpha}^{-1} = \mu_{H,\alpha^{-1}}. \quad (2.2.4)$$

This follows directly from the construction in [Marchetti(2025), Definitions 5–6].

Proof. An **automorphism** preserves proper faces and the order relation, so the moving branch remains in $\mathcal{R}(P)$. Since $\hat{\alpha}(H) = H$, the condition $F \leq H$ holds exactly when $\hat{\alpha}(F) \leq H$. Thus the moving and fixed subsets are preserved separately. On the moving subset the extension of α^{-1} is $\hat{\alpha}^{-1}$ by **uniqueness**, and cancels the first map. On the fixed subset both maps are the identity. $\square$


The general construction combining permutations of complementary subsets is [Equiv.Perm.subtypeCongr](#); its inverse formula is [Equiv.Perm.subtypeCongr.symm](#). Here the predicate is $F \leq H$, the moving permutation is the restricted shape automorphism and the other permutation is the identity. These declarations supply that general combination; the polytope-specific preservation argument is the proof above.

The inverse lemma makes the local action reversible. The tetrahedron calculation that follows is a small witness for why the facet is part of a move's identity: two locally valid turns can act differently on the same incidence record and need not commute.

Example 2.2.4 (Two turns of a tetrahedron)

Represent the **faces** of a tetrahedron by subsets of $\{1, 2, 3, 4\}$. We calculate **facet moves** from [Marchetti(2025), Definitions 5–6]. For the facet $H = 123$, take an element of the **rotation subgroup** with global extension $\hat{\alpha} = (123)$, fixing 4. The vertex-edge records follow the cycle

$$(1, 12) \mapsto (2, 23) \mapsto (3, 13) \mapsto (1, 12). \quad (2.2.5)$$



Every arrow is the local move $\mu_{H,\alpha}$. The record $(4, 14)$ stays fixed because $4 \notin H$, although the global shape symmetry sends it to $(4, 24)$. The di-

agram records this **permutation of incidence positions**, rather than the tetrahedron's face graph.

For records whose second entry is a facet, the same move gives

$$(1, 124) \mapsto (2, 234), \quad (12, 124) \mapsto (23, 234), \quad (4, 124) \mapsto (4, 124). \quad (2.2.6)$$

The last record is fixed even though the global symmetry changes its second entry. Now let $K = 124$ and take an element of the **rotation subgroup** with global extension $\widehat{\beta} = (124)$. Applying the H -move and then the K -move sends

$$(1, 124) \xrightarrow{\mu_{H,\alpha}} (2, 234) \xrightarrow{\mu_{K,\beta}} (4, 134). \quad (2.2.7)$$

Reversing their order gives $(1, 124) \mapsto (2, 124) \mapsto (3, 234)$, so the two moves do not commute. Each has order three: its cube is the identity on both branches, and its displayed action is nontrivial.

The tetrahedron witness motivates the generated move group: reachable labelled configurations are products of these reversible local turns. The ordinary cube calculation then shows how the same definition compares with visible stickers and which data it leaves out.

Definition 2.2.5 (The Rubik group and labelled configurations)

For a finite **regular abstract polytope** P of rank $n \geq 3$, the **Rubik group** is the subgroup generated by all **facet moves**:

$$\text{Rub}(P) = \langle \mu_{H,\alpha} : H \text{ a facet}, \quad \alpha \in \Gamma^+(H/F_{-1}) \rangle \leq \text{Sym}(\mathcal{R}(P)). \quad (2.2.8)$$

This is Marchetti's generated **group** of permutations [Marchetti(2025), Section 3, p. 6, following Definition 6]. The **inverse lemma** ensures that the generators are permutations.


Give each incidence position its own initial label. A **labelled configuration** can be represented by the **permutation** carrying each initial label to its current position. The identity is the solved configuration; the reachable labelled configurations are the elements of $\text{Rub}(P)$. A new move acts on the left: applying μ to configuration g gives μg , with g acting first. A visible coloring may identify several labelled configurations, so a solved visible coloring need not determine the labelled permutation.

Example 2.2.6 (One ordinary cube face turn)

Realize the cube as $[-1, 1]^3$ and let H be the facet $z = 1$. The **facet move** construction of [Marchetti(2025), Definitions 5–6] gives the following explicit outer-face action. The shape rotation

$$q(x, y, z) = (y, -x, z) \quad (2.2.9)$$

is a clockwise quarter-turn when viewed from positive z , with x rightward and y upward. Its restriction to H supplies our move. The corner $(1, 1, 1)$ goes to $(1, -1, 1)$, and its color on $y = 1$ goes to the color position on $x = 1$. The corresponding corner at $z = -1$ is fixed by the puzzle move, although the global map q moves it.



The signs give the x, y coordinates of the four corners on $z = 1$. The arrows show their clockwise motion; the central incidence record (H, H) stays fixed. No orientation of a marked center is recorded.

Restricting the second entry to facets gives $8 \cdot 3 + 12 \cdot 2 + 6 = 54$ incidence positions: three per corner, two per edge and one per face center. The turn moves twelve corner stickers and eight edge stickers, fixing the other thirty-four positions, including the unmarked center of the turned face. This is the usual outer-face action on visible stickers of a $3 \times 3 \times 3$ cube in a fixed frame.

The full incidence set has $8 \cdot 7 + 12 \cdot 3 + 6 = 98$ elements: above each vertex are that vertex, three edges and three facets; above each edge are itself and two

facets. The facet-colored subset is invariant under every move. Its restricted action supplies the comparison with ordinary stickers; a faithful restriction for all abstract polytopes is not needed. Marked center orientations and inner-layer turns are additional data, absent from this comparison.

The local action determines which positions participate and how their incident labels travel. **Frames and orientation** separate this transport into a destination and a local orientation.

3 Positions, orientations, and global constraints

After a **facet move**, a piece can reach the right position with the wrong local orientation. The first question is therefore what data records both the destination and the frame carried by the piece. **Local frames** turn this transport into group coordinates, as in [Marchetti(2025), Section 3.1]. We then derive the **composition law**, a frame-independent **total-orientation character**, and the additional test that makes this character a conserved quantity for cube face turns.

Definition 3.1 (Framed transports)

Let I be a finite set of positions, Ω a nonempty set, and K a group acting faithfully on Ω . At each position i , choose a set X_i of local labels and a bijective **local frame** $c_i : \Omega \rightarrow X_i$. Keep the fibers distinct by using the tagged union

$$E = \bigsqcup_{i \in I} X_i = \{(i, y) : i \in I, y \in X_i\}. \quad (3.1)$$

A **framed transport** is a permutation of E given by a position permutation $\pi \in \text{Sym}(I)$ and elements $a_i \in K$, with

$$g(i, c_i(x)) = (\pi(i), c_{\pi(i)}(a_i x)). \quad (3.2)$$

Write $g = (\pi, a)$. Its local factors are indexed by the source position. Nonempty fibers determine π uniquely, and faithfulness determines each a_i uniquely. Conversely every such pair gives a permutation: in framed coordinates its inverse sends (j, x) to $(\pi^{-1}(j), a_{\pi^{-1}(j)}^{-1}x)$.

This finite-position model isolates the coordinate mechanism of [Marchetti(2025), Section 3.1, Definition 7 and equation (5)]. Which local label sets and local group describe a particular puzzle must still be justified for that puzzle. The requirement that a transport moves whole fibers is part of the model.

Under composition these transports form a **wreath product** $K \wr \text{Sym}(I)$. The normal subgroup K^I acts within fibers, and position permutations reorder its factors. In `TauCeti.WreathProduct`, a pair is stored using the destination-indexed function $A_j = a_{\pi^{-1}(j)}$. Its permutation action is `TauCeti.WreathProduct.imprimitiveToPerm_apply`; the nonempty-fiber and faithful-action hypotheses give its injectivity in `TauCeti.WreathProduct.imprimitiveToPerm_injective`. The bijections c_i identify that action on $I \times \Omega$ with the action on E above.

Lemma 3.2 (Composition of framed transports)

`mul_left`)

Let I be finite and let K act faithfully on a nonempty Ω , with bijective frames $c_i : \Omega \rightarrow X_i$. Write two **framed transports** as $g = (\pi, a)$ and $h = (\sigma, b)$, acting in coordinates by $g(i, x) = (\pi(i), a_i x)$ and $h(i, x) = (\sigma(i), b_i x)$. With left actions, applying g and then h gives the **composition law**

$$hg = (\sigma\pi, (b_{\pi(i)}a_i)_{i \in I}). \quad (3.3)$$

Proof. The intermediate coordinates are $(\pi(i), a_i x)$. Applying h evaluates its local map at that intermediate position, giving $(\sigma\pi(i), b_{\pi(i)}a_i x)$. Thus a local factor must be multiplied at its transported position index. $\square$

$$\begin{array}{ccccc}
 \Omega & \xrightarrow{a_i} & \Omega & \xrightarrow{b_{\pi(i)}} & \Omega \\
 \downarrow c_i & & \downarrow c_{\pi(i)} & & \downarrow c_{\sigma\pi(i)} \\
 X_i & \xrightarrow{g} & X_{\pi(i)} & \xrightarrow{h} & X_{\sigma\pi(i)}
 \end{array}$$

The vertical arrows are the chosen frames; the bottom arrows are the restrictions of the actual transports. Each square commutes. Along the top row, the second factor is evaluated at $\pi(i)$, explaining the shifted index in the composite. This expands the wreath-product coordinates of [Marchetti(2025), Section 3.1, pp. 6–7]. The **wreath product** uses destination-indexed factors in the linked formalization.

The linked multiplication theorem uses destination-indexed factors. Set $A_j = a_{\pi^{-1}(j)}$ and $B_j = b_{\sigma^{-1}(j)}$. Its factor for hg at destination j is

$$B_j A_{\sigma^{-1}(j)} = b_{\sigma^{-1}(j)} a_{\pi^{-1}\sigma^{-1}(j)}. \quad (3.4)$$

Putting $i = (\sigma\pi)^{-1}(j)$ makes this exactly $b_{\pi(i)} a_i$. Thus the theorem and the displayed source-indexed law use the same multiplication order.

Theorem 3.3 (The abelianized total orientation)

Let I be finite, let K act faithfully on a nonempty Ω , and choose bijective frames $c_i : \Omega \rightarrow X_i$. For a **framed transport** $g = (\pi, a)$, let $K_{\text{ab}} = K/[K, K]$ be its **abelianization**, where $[K, K]$ is generated by $uvu^{-1}v^{-1}$, and write $[a]$ for the class of a . The **total-orientation character**

$$\chi(g) = \prod_{i \in I} [a_i] \quad \text{in } K_{\text{ab}} \quad (3.5)$$

is a homomorphism, independent of replacing the frames by $c'_i = c_i t_i$ for $t_i \in K$. This is the canonical character in [Marchetti(2025), Section 3.1, p. 7]; the proof below gives both assertions directly in source-indexed coordinates.

Proof. For $h = (\sigma, b)$, the **composition law** and commutativity of the target give

$$\chi(hg) = \prod_i [b_{\pi(i)}][a_i] = \left(\prod_j [b_j] \right) \left(\prod_i [a_i] \right) = \chi(h)\chi(g). \quad (3.6)$$

The middle equality reindexes using the bijection π . In the new frames the local map is $a'_i = t_{\pi(i)}^{-1} a_i t_i$. Its abelianized product contains each $[t_i]$ and its inverse exactly once, so these cancel and leave $\chi(g)$. $\square$

Frame independence here concerns changes within the chosen group of frames. Reversing a common orientation convention can also change the identification of the target group. The target and its quotient map are **Abelianization** and **Abelianization.of**; the total product and its frame independence are the argument above.

Remark 3.4 (When a character is conserved)

For finite I , the ambient coordinate group $W = K \wr \text{Sym}(I)$ has the **total-orientation character** $\chi(\pi, a) = \prod_i [a_i] \in K_{\text{ab}}$. If a puzzle's allowed moves form a set $S \subseteq W$, its legal group is $G = \langle S \rangle$, usually a proper subgroup of W . Then

$$G \subseteq \ker \chi \iff \chi(s) = 1 \text{ for every } s \in S. \quad (3.7)$$

The forward implication includes each generator. Conversely the kernel of a homomorphism is a subgroup, so it contains all products and inverses of generators lying in it. This is the generator-check step discussed after the canonical character in [Marchetti(2025), Section 3.1, p. 7]. Its premise must be verified for the chosen moves.

The nonabelian product of local factors need not be independent of order or frames. Abelianization makes reindexing and frame cancellation valid, but neither of those operations says that $\chi(s) = 1$. The **cycle criterion** supplies a sufficient test, applied explicitly to **cube face turns**.

Lemma 3.5 (Orientation checked cycle by cycle)

Let $g = (\pi, a)$ be a **framed transport** on finitely many positions I , with a faithful local group K on nonempty Ω and bijective local frames. On a cycle

$$i_0 \mapsto i_1 \mapsto \cdots \mapsto i_{r-1} \mapsto i_0 \quad (3.8)$$

of π , the return transport in the first fiber has framed map $a_{i_{r-1}} \cdots a_{i_1} a_{i_0}$. Its class in K_{ab} is that cycle's contribution to $\chi(g) = \prod_{i \in I} [a_i]$. Consequently, identity return transport on every position cycle implies $\chi(g) = 1$. This is the **cycle criterion**.

Proof. Apply the coordinate action r times. Each local map is applied after the preceding one, giving the stated product with the first map on the right. In the abelianization its factors may be reordered; multiplying over disjoint cycles uses each factor exactly once. A fixed position is a cycle of length one and must also be checked. $\square$

This is a sufficient test derived from the canonical character of [Marchetti(2025), Section 3.1, p. 7]. A generator's finite order alone does not make its return map on each shorter position cycle the identity; the local transport on that cycle is the quantity to check.

Example 3.6 (Corner twists and edge flips on the cube)

Use the unmarked-center, fixed-frame sticker model of **one cube face turn**. We apply the canonical character of [Marchetti(2025), Section 3.1, p. 7] to this explicit action. At a corner the three incident sticker directions have a cyclic order inherited from the ambient orientation. Choose bijective cyclic frames at all eight positions. Rotations preserve that order, giving local group C_3 ; at each of the twelve edges an ordering of its two stickers gives local group C_2 . Both actions on the local stickers are faithful. Writing these groups additively, the two characters are

$$\chi_c = \sum_{i=1}^8 a_i \pmod{3}, \quad \chi_e = \sum_{j=1}^{12} b_j \pmod{2}. \quad (3.9)$$

A quarter-turn has one four-cycle of affected corners and one of affected edges. Four applications restore every sticker, so the return transport on each of these cycles is the identity. Every other corner and edge is fixed with identity local transport. The **cycle criterion** gives $\chi_c = 0$ and $\chi_e = 0$ for every face quarter-turn. The **kernel argument** extends this to every word in those turns, including inverses and half-turns.

For a concrete cyclic frame, choose any frame at corner i_0 of the upper face and define the next three by transporting it with the shape rotation q . Since $q^4 = 1$, all four local factors of this turn become zero. This frame checks that one generator; the **frame-independence theorem** transfers the conclusion to any fixed choice of cyclic frames used for the other generators.

A configuration with a **corner twist** of just one corner by 1 has $\chi_c = 1$ and is impossible under these moves. Twisting two corners by 1 and -1 passes this test, but the calculation alone does not construct a legal sequence. Likewise, flipping exactly one edge fails the second test. An **edge flip** of exactly one edge therefore fails the second test. Reachability needs additional arguments.

Transfer gives a group-theoretic form of the same frame cancellation and shows how the local group limits what total orientation can detect.

4 Transfer and the effect of dimension

The **total-orientation character** detects one conservation law, but it does not explain how to compare frames chosen at different positions. The motivat-

ing question is how to transport that comparison through a transitive action. **Transfer** uses the same cancellation as the **total-orientation character** for coset representatives [Mayer(2015), Sections 3.1–3.3]. The first section identifies the connection through an actual transitive action and its **stabilizer**; the second computes how the **local rotation groups** of simplices and cubes change with dimension, following [Marchetti(2025), Sections 4–5].

4.1 Transfer and local orientation

Left multiplication permutes the cosets of a finite-index subgroup. Representatives record the motion between cosets, leaving one residual subgroup element at each position. Abelianizing their product gives **transfer** [Mayer(2015), Sections 3.1–3.3]. Transported frames then identify its local image with the **total-orientation character** of an action.

Theorem 4.1.1 (Finite-index transfer)

`MonoidHom.transfer` `MonoidHom.transfer_def` `transfer_eq_prod_of_bijective`)

Let $H \leq G$ have finite index. Choose representatives r_i for the left cosets $r_i H$, indexed bijectively by a finite set I . For $g \in G$, left multiplication permutes the cosets, so uniquely write

$$g r_i = r_{\pi_g(i)} h_i(g), \quad h_i(g) \in H. \quad (4.1.1)$$

Write $H_{\text{ab}} = H/[H, H]$ and $[h]$ for the class of $h \in H$. The **abelianization** is the target of the **transfer**

$$T_{G,H} : G \longrightarrow H_{\text{ab}}, \quad T_{G,H}(g) = \prod_{i \in I} [h_i(g)] \quad (4.1.2)$$

is a homomorphism independent of the chosen representatives. This is the left-coset construction of [Mayer(2015), Definition 3.1, Proposition 3.1 and Theorem 3.1, pp. 4–6]. Neither finiteness of G nor normality of H is required.

Proof. We compose left actions rightmost first. Multiplication on cosets gives $\pi_{gk} = \pi_g \pi_k$. Expanding gkr_i in two steps yields

$$h_i(gk) = h_{\pi_k(i)}(g) h_i(k). \quad (4.1.3)$$

Taking classes and multiplying over i gives $T_{G,H}(gk) = T_{G,H}(g)T_{G,H}(k)$, because π_k reorders the first factors. Replace r_i by $r_i t_i$, with $t_i \in H$. The new residual factor is $t_{\pi_g(i)}^{-1} h_i(g) t_i$. Its total class contains each $[t_i]$ and its inverse exactly once, so they cancel. Reindexing the cosets also only reorders an abelian product. This proves independence of all choices. $\square$

The linked transfer definitions specialize the homomorphism $\phi : H \rightarrow A$ to the quotient $H \rightarrow H_{\text{ab}}$. The indexed formula takes $f(i) = r_i$ and $\pi = \pi_g$; its factor $r_{\pi_g(i)}^{-1} g r_i$ is precisely $h_i(g)$. Its finite-index instance follows from the finite bijective family of cosets.

Proposition 4.1.2 (Transfer computes total local orientation)

Let G act transitively on a finite nonempty position set I , and on a tagged union of fibers X_i , carrying X_i bijectively to X_{gi} . Fix $i_0 \in I$ and put $H = G_{i_0}$. Let K act faithfully on a nonempty set Ω ; choose a bijective base frame $c_0 : \Omega \rightarrow X_{i_0}$ and a homomorphism $\rho : H \rightarrow K$ satisfying $h c_0 = c_0 \rho(h)$ for every $h \in H$. Choose $r_i \in G$ with $r_i i_0 = i$, and set $c_i = r_i c_0$. Then the **total-orientation character** of this **framed transport** is

$$\chi = \rho_{\text{ab}} \circ T_{G,H} : G \longrightarrow K_{\text{ab}}, \quad (4.1.4)$$

where $T_{G,H}$ is **transfer** and $\rho_{\text{ab}}([h]) = [\rho(h)]$. This identifies the coset construction in [Mayer(2015), Section 3.3, p. 6] with the local-frame coordinates of [Marchetti(2025), Section 3.1, pp. 6–7].

Proof. The map $G/H \rightarrow I$, $rH \mapsto r i_0$, is a bijection: $r i_0 = s i_0$ exactly when $s^{-1}r \in H$. Thus H has finite index $|I|$, and the r_i are left-coset representatives. Write $\pi_g(i) = gi$ and $g r_i = r_{\pi_g(i)} h_i(g)$. Then

$$g c_i = g r_i c_0 = r_{\pi_g(i)} h_i(g) c_0 = c_{\pi_g(i)} \rho(h_i(g)). \quad (4.1.5)$$

Consequently every local factor belongs to K , and its source-indexed coordinate is $a_i(g) = \rho(h_i(g))$. A homomorphism sends commutators to commutators, so ρ_{ab} is well defined. Multiplying the local classes gives

$$\chi(g) = \prod_i [\rho(h_i(g))] = \rho_{\text{ab}} \left(\prod_i [h_i(g)] \right). \quad (4.1.6)$$

$\square$

$$\begin{array}{ccc}
\Omega & \xrightarrow{\rho(h_i(g))} & \Omega \\
\downarrow r_i c_0 & & \downarrow r_{\pi_g(i)} c_0 \\
X_i & \xrightarrow{g} & X_{\pi_g(i)}
\end{array}$$

The vertical arrows are the transported frames. The equality $gr_i = r_{\pi_g(i)}h_i(g)$ makes the square commute: after the destination representative is removed, the remaining motion fixes i_0 and acts through ρ . Changing representatives changes these frames within K , so the character is unchanged.

The induced map is [Abelianization.map](#). Naturality of transfer is [MonoidHom.transfer_comp](#), with $\phi : H \rightarrow H_{\text{ab}}$ the quotient map and the target map ρ_{ab} . That algebraic identity supports the last equality; the identification with the actual fiber action is the transported-frame argument above.

Example 4.1.3 (Corner and edge stabilizers of the cube)

Let G be the actual permutation group generated by the six outer-face quarter-turns of the ordinary cube, with its frame of reference fixed. Use the **corner and edge fibers**: cyclic corner labels with local group C_3 , and edge labels with local group C_2 . This specializes the local-orientation construction of [\[Marchetti\(2025\), Section 3.1\]](#); the orbit argument below identifies the stabilizers needed for **transfer**.

The action on eight corner positions is transitive. Adjacent corners share a face whose quarter-turn takes one to the other, and the cube's edge graph is connected. Hence a **corner stabilizer** H_c has index eight. Its action on the three cyclically ordered labels gives a homomorphism $\rho_c : H_c \rightarrow C_3$. The action on twelve edge positions is also transitive: the four edges of one face form a single orbit under its quarter-turn; neighboring faces share an edge; and the face-adjacency graph is connected. Thus an **edge stabilizer** H_e has index twelve, with local action $\rho_e : H_e \rightarrow C_2$.

Choose a frame at the fixed corner or edge and transport it using coset

representatives. The transfer comparison gives

$$\chi_c = (\rho_c)_{\text{ab}} \circ T_{G, H_c}, \quad \chi_e = (\rho_e)_{\text{ab}} \circ T_{G, H_e}. \quad (4.1.7)$$

We identify the abelianizations of C_3 and C_2 with those groups. The **face-turn calculation** proves that the twist and flip characters vanish on every legal generator, hence on G . It follows that these two local images of transfer vanish. It does not follow that T_{G, H_c} or T_{G, H_e} is trivial in the larger stabilizer abelianizations. No classification of G was needed.

4.2 Dimension and local rotation groups

The **local corner groups** of simplices and cubes are **alternating groups**. Their **commutator subgroups** determine the information that survives the **total-orientation character** [Marchetti(2025), Sections 4–5].

Proposition 4.2.1 (Local rotations of simplices and cubes)

For $d \geq 3$, a **vertex figure** of a d -simplex or a d -cube is a $(d - 1)$ -simplex. Its automorphism group is S_d , and its **rotation subgroup** is A_d . These are the **local corner groups** used in [Marchetti(2025), Section 4, p. 10, and Section 5, p. 14].

Proof. At a vertex of either polytope there are d incident edges. Faces containing that vertex correspond to subsets of those edges: for a simplex one chooses the other vertices, and for a cube one chooses the coordinate directions in which the face varies. This is the face poset of a simplex on d vertices, of dimension $d - 1$. An automorphism permutes its vertices, and every vertex permutation extends uniquely to the subsets, giving S_d .

A base flag orders these d vertices. Its distinguished involutions are adjacent transpositions. They generate S_d , and the even words are precisely the even permutations. Thus the rotation subgroup is A_d . $\square$

Here d is abstract dimension. Varying the layer count of an $N \times N \times N$ cube leaves the ambient dimension equal to three; it changes the piece orbits and allowed layer moves instead.

Theorem 4.2.2 (Abelianizations of the local alternating groups)

Let A_d be the **alternating group** on d letters, with composition rightmost first. Then

$$(A_3)_{\text{ab}} \cong C_3, \quad (A_4)_{\text{ab}} \cong C_3, \quad (A_d)_{\text{ab}} = 1 \quad (d \geq 5). \quad (4.2.1)$$

The corresponding **commutator subgroups** are stated after Proposition 4.1 in [Marchetti(2025), Section 4.1, p. 11]. We give the elementary calculation, using $[g, h] = ghg^{-1}h^{-1}$.

Proof. The group A_3 is cyclic of order three, so already abelian. In A_4 , let $V = \{1, (12)(34), (13)(24), (14)(23)\}$. Products and inverses stay in this four-element set, and conjugation permutes its three nonidentity elements. Hence V is normal. The quotient has order three, so it is cyclic and $[A_4, A_4] \subseteq V$. Direct multiplication gives $[(123), (124)] = (12)(34)$. Conjugating this element by the three-cycle (123) gives the other two double transpositions. The commutator subgroup is normal, so it contains all of V ; equality follows.

For $d \geq 5$, put $g = (12)(34)$ and $h = (12)(35)$ in A_d . Multiplication gives $[g, h] = (345)$. Every three-cycle is conjugate to this one within A_d : first choose a conjugating permutation in S_d ; if it is odd, compose on the right with a transposition of two symbols outside $\{3, 4, 5\}$. This centralizes (345) and changes the parity. Thus the normal subgroup $[A_d, A_d]$ contains every three-cycle.

Three-cycles generate all even permutations. Write an even permutation as a product of an even number of transpositions and pair them. Equal transpositions cancel; two sharing one symbol give a three-cycle; a disjoint pair satisfies $(ab)(cd) = (acb)(acd)$. Hence every pair is a product of three-cycles, proving $[A_d, A_d] = A_d$. $\square$

The three cases have separate formalized ingredients. `alternatingGroup.isCyclic_of_card_le_three` gives the cyclicity of A_3 . For four letters, `alternatingGroup.kleinFour_eq_commutator` identifies V with the commutator subgroup, and `TauCeti.card_abelianization_alternatingGroup` gives the quotient's cardinality three; the displayed cyclic isomorphism then uses the elementary classification of groups of prime order. For $d \geq 5$, `commutator_alternatingGroup_eq_top` is the perfection statement. These links concern alternating groups, not the entire puzzle group.

Remark 4.2.3 (What the corner character can detect)

For a d -simplex or d -cube with $d \geq 3$, the **local corner groups** have rotation subgroup A_d . Their contribution to the **total-orientation character** passes through $(A_d)_{\text{ab}}$. The **abelianization calculation** therefore measures exactly what this detector can see, as in [Marchetti(2025), Section 4.1, p. 11, and equation (19), p. 14].

In dimension three the target is C_3 , the familiar corner twist sum. In dimension four the local group has grown to the nonabelian A_4 , but the character still sees only a C_3 quotient: every local change in the Klein four group disappears from it. In dimensions five and above the local group is perfect, so this corner character is identically trivial on the ambient framed transports.

A lost detector does not classify the legal moves. Other piece types, permutation signs and the subgroup actually generated by facet turns may still impose constraints. Even when the target is nontrivial, **conservation** requires a generator check; the local abelianization alone does not supply one.

References

- [Izmestiev and Joswig(2002)] Ivan Izmestiev and Michael Joswig. 2002. Branched Coverings, Triangulations, and 3-Manifolds. arXiv:math/0108202 <https://arxiv.org/abs/math/0108202v2> (Cited in section (document))
- [Marchetti(2025)] Giovanni Luca Marchetti. 2025. Rubik’s Abstract Polytopes. arXiv:2502.13518 [math.CO] <https://arxiv.org/abs/2502.13518v1> (Cited in sections (document), 2, 2.2, 2.2.3, 2.2.4, 2.2.5, 2.2.6, 3, 3.1, 3.2, 3.3, 3.4, 3.5, 3.6, 4, 4.1.2, 4.1.3, 4.2, 4.2.1, 4.2.2, and 4.2.3)
- [Mayer(2015)] Daniel C. Mayer. 2015. Artin Transfer Patterns on Descendant Trees of Finite p-Groups. arXiv:1511.07819 [math.GR] <https://arxiv.org/abs/1511.07819v1> (Cited in sections 4, 4.1, 4.1.1, and 4.1.2)
- [Pan and Kondor(2021)] Horace Pan and Risi Kondor. 2021. Fourier Bases for Solving Permutation Puzzles. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 130)*. PMLR, 172–180. <https://proceedings.mlr.press/v130/pan21a.html> (Cited in section (document))
- [Schulte(2017)] Egon Schulte. 2017. Regular Incidence Complexes, Polytopes, and C-Groups. arXiv:1711.02297 [math.CO] <https://arxiv.org/abs/1711.02297v1> (Cited in sections 2, 2.1, (document), 2.1.3, 2.1.6, 2.1.6, and 2.1.7)

Alphabetical Index

- Abelianization, 18
- Abstract n-polytope, 7
- Alternating group, 25
- Automorphism, 9
- Commutator
 - subgroups, 25
- Composition law, 17
- Corner stabilizer, 23
- Corner twist, 20
- Cycle criterion, 19
- Diamond condition, 7
- Distinguished
 - generators, 10
- Edge flip, 20
- Edge stabilizer, 23
- Faces, 7
- Facet move, 12
- Flag, 7
- Framed transport, 16
- Incidence positions, 12
- Intersection property, 10
- Labelled configuration, 14
- Local corner groups, 24
- Local frame, 16
- Permutation, 12
- Ranks, 7
- Regular, 9
- Rotation subgroup, 11
- Rubik group, 14
- Section, 8
- String C-group, 10
- Strong flag connectivity, 7
- Total-orientation
 - character, 18
- Transfer, 21
- Vertex figure, 8
- Wreath product, 17