Duality as an Incidence Category Picture: A Categorical Rendering of the Cube–Octahedron Pair

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

December 28, 2025

Polyhedral duality is usually taught as a geometric trick: take face-centers, connect them, and out drops the dual. Category theory offers a different lens. Instead of treating a polyhedron as a solid object in space, we treat it as a structured family of cells—vertices, edges, faces—tied together by incidence. That incidence data naturally forms a small category: objects are the cells, and morphisms encode "is incident to" (vertex to edge, edge to face), with composition capturing higher incidence (vertex to face). In this view, duality is not merely a spatial construction but a contravariant transformation: a reversal of incidence that swaps dimensions, sending faces to vertices and (in three dimensions) vertices to faces while preserving edges in a precise sense. This paper develops that perspective using a concrete, page-sized "category picture" of the cube and its dual octahedron. The figure serves as the central artifact: two incidence categories laid out in parallel strata, with an explicit duality map drawn between cube faces and octahedral vertices. From this starting point we formalize the incidence category, describe the duality functorial behavior as an opposite-category phenomenon, and show how the same method generalizes to other Platonic pairs and beyond. The goal is a usable blueprint: a rigorous yet visual way to see duality as structure-preserving reversal, not magic.

Keywords: polyhedral duality, incidence category, face poset, opposite category, contravariant equivalence, cube, octahedron, Platonic solids, functoriality, Hasse diagram, combinatorial topology, Schlegel diagram.

A collaboration with GPT-5.2-Thinking. 66 pages.

Outline

1. The Figure as Thesis: What a "Category Picture" Shows

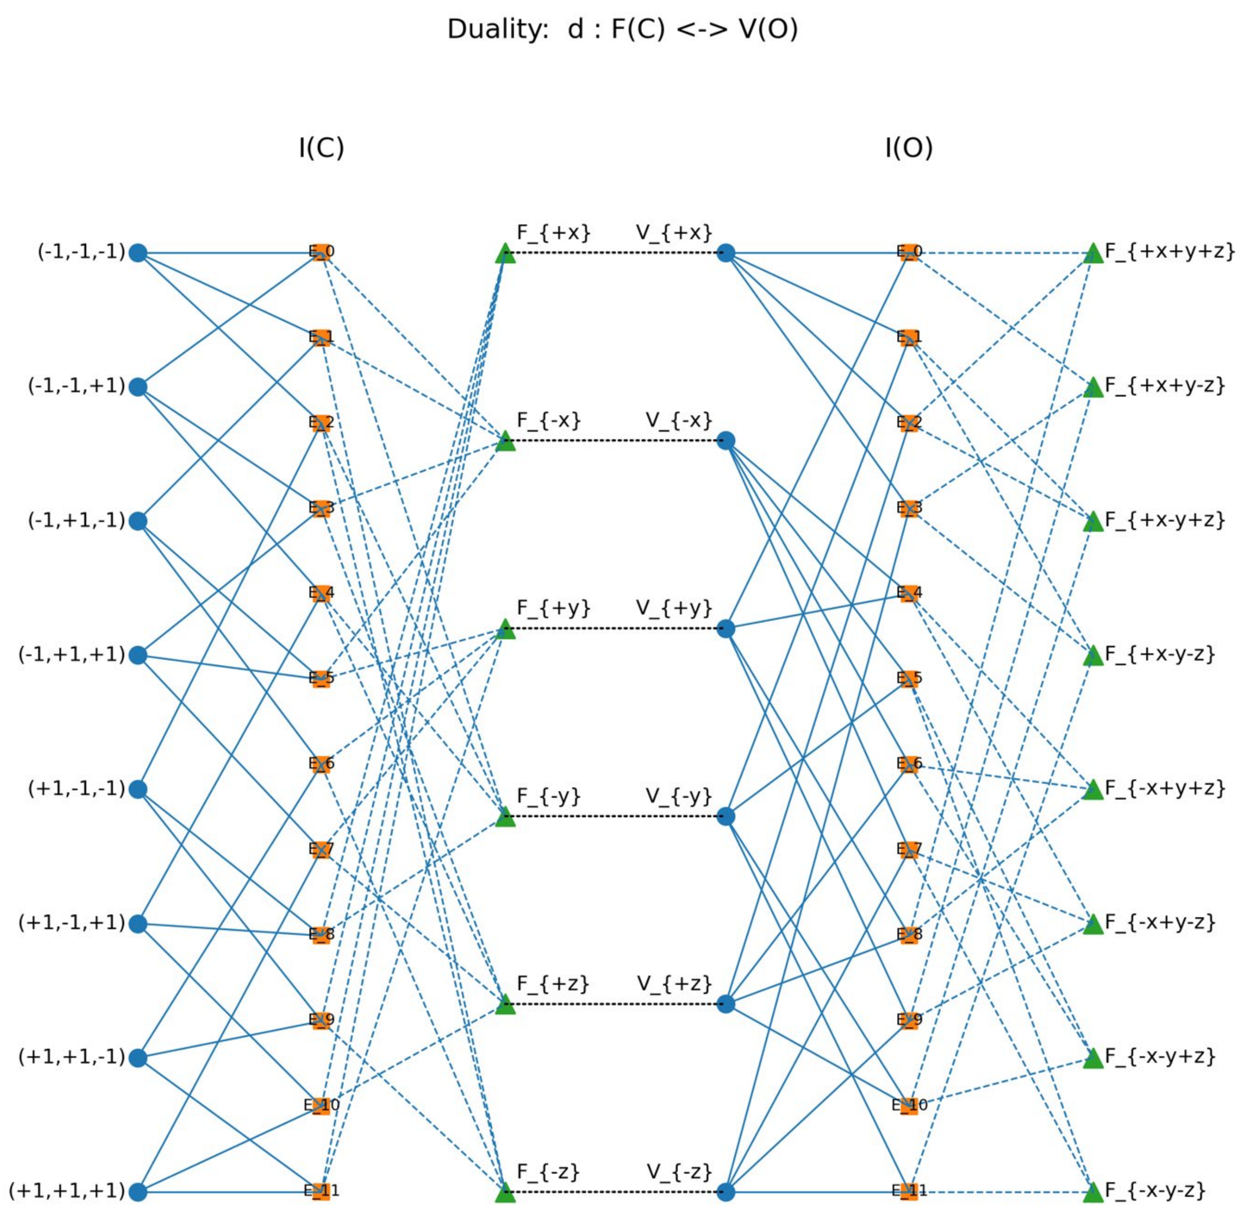

Figure 1. The figure is not meant to be a decorative diagram of the cube and octahedron as shapes in space. It is a compact claim about what the cube and octahedron *are* when we stop thinking of them as solids and instead treat them as structured data. In this paper, the governing idea is that a polyhedron can be replaced by its incidence structure, and that incidence structure can be organized as a category. Once you do that, "duality" stops being a geometric trick and becomes a contravariant relationship between two categories: a systematic reversal of how incidence flows across dimensions. The figure is the thesis rendered visually: two incidence categories presented in parallel, plus a partial duality correspondence drawn between them.

1.1 Two panels, one claim: incidence as a category

The left panel, labeled I(C), is intended to be read as a small category built from the cube C. The right panel, labeled I(O), is the analogous category for the octahedron O. The figure asserts that for each polyhedron P there is a canonical “incidence category” I(P) whose objects are the polyhedron’s cells by dimension (in 3D: vertices, edges, faces) and whose morphisms encode incidence relations. In ordinary geometric language, the data “vertex lies on edge” and “edge bounds face” is an adjacency story. In categorical language, the same story is expressed as arrows and their composites: if v -> e and e -> f, then there is a composite v -> f. That composite is not an extra geometric primitive; it is forced by the categorical closure under composition, and it corresponds to the derived incidence fact that the vertex v lies on the face f.

The figure’s two-panel layout makes a second claim: the cube and octahedron, as a dual pair, carry incidence structures that mirror each other under dimension reversal. The reader does not need to “see” a cube or octahedron in the Euclidean sense. The figure is already the cube and octahedron at the level that duality cares about: the combinatorics of which cells touch which.

1.2 Reading conventions: objects by dimension and arrows by incidence

Each panel is arranged in three vertical columns that correspond to the three dimensions of cells in a convex polyhedron.

In I(C) on the left:

- The leftmost column shows the cube’s vertices V(C) as blue dots, labeled by the coordinate model (+/-1, +/-1, +/-1). This labeling is not decorative: it fixes a concrete combinatorial model for the cube and makes the edge relation “differ in exactly one coordinate” literally readable.
- The middle column shows the cube’s edges E(C) as orange squares, labeled E_0, E_1, ..., E_11. The edge labels are just names; what matters is their incidence to vertices and faces.
- The rightmost column shows the cube’s faces F(C) as green triangles, labeled F_{+x}, F_{-x}, F_{+y}, F_{-y}, F_{+z}, F_{-z}. These are the six coordinate faces of the cube.

In I(O) on the right:

- The leftmost column shows the octahedron’s vertices V(O) as blue dots, labeled V_{+x}, V_{-x}, V_{+y}, V_{-y}, V_{+z}, V_{-z}. These correspond to the six axis directions, matching the idea that the octahedron is the convex hull of the coordinate unit vectors and their negatives (up to scale).

- The middle column shows the octahedron’s edges E(O) as orange squares, again labeled E_i.
- The rightmost column shows the octahedron’s faces F(O) as green triangles, labeled by sign triples such as $F_{+x +y +z}$, $F_{+x +y -z}$, and so on, for a total of eight.

Within each panel, the lines have the meaning “there exists an incidence morphism.” The visual style distinguishes the two primitive incidence types:

- Solid blue lines correspond to morphisms of the form $v \to e$, meaning the vertex v is incident to the edge e.
- Dashed blue lines correspond to morphisms of the form $e \to f$, meaning the edge e is incident to the face f (equivalently, e lies in the boundary of f).

The figure does not explicitly draw the composites $v \to f$, but they are present implicitly. Every time you can trace a path from a vertex to an edge to a face, the category forces a composite morphism. This is the main advantage of the categorical packaging: it makes “derived incidence” into a formal consequence of composition rather than something you must separately define.

Across the center gap there are dotted black lines, and they play a different role. They are not incidence morphisms inside either category. They are correspondences between *the two* categories, specifically a partial display of the duality map labeled in the title:

- $d : F(C) \leftrightarrow V(O)$.

Each dotted line identifies a cube face $F_{\pm x}$, $F_{\pm y}$, $F_{\pm z}$ with the corresponding octa vertex $V_{\pm x}$, $V_{\pm y}$, $V_{\pm z}$. This is one leg of the duality: faces of the cube correspond to vertices of the octahedron.

The figure does not draw the other two legs, but it strongly suggests them:

- $V(C) \leftrightarrow F(O)$ (8 cube vertices correspond to 8 octa faces),
- $E(C) \leftrightarrow E(O)$ (12 edges correspond to 12 edges).

Those correspondences are part of the full duality structure, and later sections will describe them formally even if they are not drawn in the central gap.

1.3 What is proved visually vs what must be proved formally

The figure can “prove” some things in the sense of making them inspectable at a glance, but it cannot replace formal verification. It is best understood as a structured witness: it shows you exactly what must be true if the categorical story is correct, and it makes it easy to spot contradictions.

What the figure gives you essentially for free:

- It exhibits the correct cardinalities: $|V(C)| = 8$, $|E(C)| = 12$, $|F(C)| = 6$; and $|V(O)| = 6$, $|E(O)| = 12$, $|F(O)| = 8$. Duality's dimension swap is reflected in the 6↔6 and 8↔8 pattern.
- It displays incidence patterns: each cube edge touches exactly 2 vertices and lies on exactly 2 faces; each cube face has 4 vertices and 4 edges; each octa face has 3 vertices and 3 edges; each octa vertex touches 4 edges, and so on. These degree facts are not abstract; they are encoded by the number of lines hitting each node.
- It makes the face-to-vertex duality d : F(C) -> V(O) explicit. The six dotted correspondences make it impossible to "handwave" which face maps to which vertex. The map is pinned down.

What the figure does not prove by itself, and what must be established formally in the text:

- That the construction I(P) is well-defined as a category with the intended properties. Concretely, we must specify the objects, define the generating morphisms, and state precisely what equalities (if any) we impose among composites. There are at least two natural variants: a thin category (poset-style) or a free category on the incidence graph modulo some relations. The figure is compatible with either, but the paper must choose and justify one.
- That the duality correspondence extends to an actual contravariant equivalence between incidence categories (or between the relevant subcategories / face-poset categories). The dotted lines show only one component of the duality on objects. A categorical equivalence requires a mapping on morphisms that respects incidence in the appropriate reversed direction, and it requires the usual properties: essentially surjective on objects and fully faithful (or the poset-category equivalent condition).
- That the correspondences respect composition. Informally, the duality must carry a chain of incidences in one polyhedron to a reversed chain in the dual. In symbols, if incidence yields v -> e -> f in I(C), then under the duality you should obtain something like d(f) -> d(e) -> d(v) in the opposite direction in I(O), once we define what d does to edges and vertices and in what direction it acts.
- That the chosen labeling scheme matches the geometric duality and not merely an arbitrary relabeling. The figure suggests a geometrically meaningful identification (coordinate faces ↔ coordinate vertices), but this must be stated as a theorem about the cube and octahedron models we are using.

So the figure's role is to concentrate the entire argument into a single page: it is the "interface" that makes the categorical claims legible. The paper's job is then to formalize exactly what the picture is claiming, to state the functorial duality precisely (usually as an equivalence to an opposite category), and to show that the cube–octahedron case is not a one-off curiosity but the simplest instance of a general method.

When you're ready, prompt me to expand **Section 2** and I'll keep the same house rules (no in-line references, bold headings only).

2. From Polyhedron to Data: Cells, Incidence, and Closure Under Composition

A polyhedron can be treated as geometry, but it can also be treated as a finite, structured dataset. The dataset consists of "cells" (vertices, edges, faces) together with relations saying which cells touch which. This shift is not a loss of information for many purposes. For duality in particular, the essential structure is combinatorial: duality is determined by incidence, not by metric length or angle. The key move of this paper is to package that incidence data as a category. The category then becomes the object of study, and the familiar geometric facts become categorical consequences of how the generating arrows compose.

2.1 The cell sets V, E, F and the incidence relations

Fix a convex polyhedron P in 3D (later we generalize). At the level we need for the figure, we take its three strata of cells:

- V(P) the set of vertices (0-cells),
- E(P) the set of edges (1-cells),
- F(P) the set of faces (2-cells).

For the cube C, $|V(C)|=8$, $|E(C)|=12$, $|F(C)|=6$. For the octahedron O, $|V(O)|=6$, $|E(O)|=12$, $|F(O)|=8$.

To form incidence data, we do not begin with all possible relations at once. We begin with the primitive ones that are "one step apart in dimension":

- A vertex-edge incidence relation $I_0 \subseteq V(P) \times E(P)$,
- An edge-face incidence relation $I_1 \subseteq E(P) \times F(P)$.

Intuitively:

- (v,e) in I_0 means "vertex v is an endpoint of edge e,"

- (e,f) in I_1 means “edge e lies on the boundary of face f.”

These relations can be specified in many equivalent ways depending on representation. In a combinatorial model, one might specify each edge as an unordered pair of vertices and each face as a cycle of edges. In a geometric model, one might test whether a vertex lies on a segment, or whether a segment lies in a planar polygon. For this paper, we treat incidence as given data, whether computed or supplied.

The important point is that I_0 and I_1 are enough to generate all the incidence facts we typically talk about. “Vertex lies on face” is not primitive in this approach; it is derived. Once you accept that there are vertices, edges, and faces, and you know which vertices lie on which edges and which edges lie on which faces, you can deduce whether a vertex lies on a face by passing through an edge.

This suggests the categorical step: view incidence as a directed relation that naturally composes.

2.2 Why composition encodes “vertex lies on face”

Category theory enters as soon as we say: relations should be closed under composition in the sense that if one incidence follows another, then there is a derived incidence.

In categorical language, we build a category I(P) whose objects are the cells:

- Ob(I(P)) = V(P) ⊔ E(P) ⊔ F(P).

We then add generating morphisms corresponding to the primitive incidence relations:

- for each (v,e) in I_0, a morphism v -> e,
- for each (e,f) in I_1, a morphism e -> f.

Now the category axioms force a new phenomenon that matches geometric intuition: the closure under composition. If we have

- v -> e and e -> f,

then we must have a composite morphism

- v -> f.

This composite morphism encodes the statement “v lies on f” in the sense that there exists an edge incident to both. In other words, “vertex lies on face” becomes the existence of a length-2 path in the incidence graph, and in the category it becomes the existence of a composite arrow.

This is more than convenient bookkeeping. It changes the logic of the construction. In a non-categorical treatment, you might define "vertex lies on face" as a separate relation $I\_2 \subseteq V(P) \times F(P)$, and then you would need to prove consistency conditions such as:

- if (v,e) in I_0 and (e,f) in I_1 then (v,f) in I_2.

But in the categorical treatment, you do not need to separately define I_2 at all. The condition is replaced by a structural axiom: composition exists. Derived incidence is no longer a definition you must remember; it is something the framework enforces.

This also clarifies why the arrows in the figure run from lower dimension to higher dimension. The direction is not arbitrary. It is chosen so that composition has the intended meaning: a chain V -> E -> F is a path "up the dimension ladder," and its composite is precisely the derived relationship between endpoints.

One can push this one step further conceptually. Once you admit composites, you also implicitly admit that there may be multiple ways to go from v to f: different edges on the same face might witness the same vertex-face incidence. That raises a question. Should different witnessing paths produce the same morphism v -> f, or should they be regarded as distinct morphisms? The answer depends on what category you want, and it leads directly to the thin vs non-thin distinction.

2.3 Thin categories and when incidence becomes a poset

There are two natural categorical choices for packaging incidence, both legitimate, and the figure is compatible with either. Choosing between them is not a matter of taste; it is a matter of what you intend the category to remember.

The thin (poset-like) version is the simplest and is often the most appropriate for polyhedral duality.

In a thin category, between any two objects there is at most one morphism. The category is essentially a preorder (and often a partial order). For polyhedra, the most standard "thin" construction is the face poset (or face lattice when suitable). The underlying relation is "is a face of" or "is contained in," including all dimensions. In that setting:

- we typically say there is a morphism a -> b iff the cell a is contained in the closure of b.

If we adopt this convention, then any two incidence paths witnessing v -> f collapse to the same unique morphism, because the only question is whether v is incident to f, not how many ways it is incident. The thin incidence category is therefore a categorical encoding of the yes/no incidence structure. It is exactly what you want when you care about duality as order reversal, because duality is fundamentally about reversing containment.

When does this incidence category become a poset rather than merely a preorder? It becomes a poset when antisymmetry holds: if there are morphisms a -> b and b -> a, then a=b. For a well-formed polyhedron with cells ordered by inclusion, antisymmetry holds because a proper face cannot contain a strictly larger face. So the face relation is a partial order, and the corresponding category is thin and skeletal.

The non-thin version keeps track of witnesses. One way to get it is to start with the bipartite incidence graph connecting V to E and E to F, form the free category on that directed graph, and then impose only the relations you truly want. In that category:

- a morphism is literally a path of incidences,
- different paths from v to f remain distinct unless you explicitly identify them.

This richer structure remembers multiplicity and routing information. It can be useful if your goal is to treat incidence paths as combinatorial objects in their own right, or if you want to move toward higher-categorical ideas where different paths might be related by “2-morphisms” expressing homotopies or coherence. But it is usually more information than polyhedral duality strictly needs, because duality is already determined at the level of the face poset.

For the purposes of the figure and the main claims of this paper, the thin view is the intended baseline: we care about the structured pattern of incidence, and we want duality to appear as an “op” phenomenon, which is most naturally stated for posets and thin categories. That said, the non-thin view shadows the thin one, and it is conceptually valuable because it explains why category theory feels native here. The category axioms are not an external imposition. They are exactly what you get when you treat “incidence steps” as composable processes and ask for closure.

So the paper’s stance is:

- start with the cell sets V, E, F,
- specify the primitive incidence relations V -> E and E -> F,
- enforce closure under composition so that V -> F is derived,
- and, for duality, pass to the thin quotient where incidence is a yes/no relation and the structure becomes a poset-category.

This is the bridge from “polyhedron as solid” to “polyhedron as categorical object.” The remaining sections will show that once the bridge is crossed, duality is no longer mysterious: it becomes a dimension-reversing correspondence that reads naturally as a contravariant equivalence, with the cube–octahedron diagram serving as the simplest explicit model.

3. Defining the Incidence Category I(P)

The phrase "incidence category" can be made precise in more than one way. The figure suggests a particular, pragmatic definition: objects are cells, arrows are generated by incidence, and composition expresses derived incidence. But there is a subtle design choice lurking underneath: do we want morphisms to remember the *path* by which incidence is witnessed, or do we want morphisms to collapse to a simple yes/no relation ("is incident to")? The first choice leads naturally to a free category built from an incidence graph. The second choice leads to a thin category (a poset-category) that is closer to the traditional face poset. Both are valid, and both can be connected by a canonical quotient. In this section we define I(P) in a way that makes the tradeoff explicit and keeps the duality story clean.

3.1 Objects as cells and generating morphisms

Let P be a convex polyhedron (or more generally, a regular cell complex in dimension 3). We start with the sets of cells by dimension:

- V(P) vertices,
- E(P) edges,
- F(P) faces.

We then declare the object set of the incidence category to be the disjoint union:

- Ob(I(P)) = V(P) ⊔ E(P) ⊔ F(P).

Now we need morphisms. The basic incidence data is:

- vertex-edge incidence: for each edge e, there are exactly two endpoints, giving pairs (v,e) with v in endpoints(e);
- edge-face incidence: for each face f, the boundary is a cycle of edges, giving pairs (e,f) with e in boundary(f).

These give the generating morphisms:

- for each incident pair (v,e), include a generating morphism v -> e;
- for each incident pair (e,f), include a generating morphism e -> f.

So far, this is just a directed graph description: objects are nodes and generating morphisms are directed edges. The categorical leap is that we insist on identities and composition. In particular:

- each object x has an identity morphism id_x : x -> x;

- whenever a -> b and b -> c are morphisms, their composite a -> c exists and is associative.

At this point, you already get an immediate payoff: any time v lies on an edge e and that edge lies on a face f, there exists a composite morphism v -> f. Thus the category automatically contains derived vertex-face incidence.

However, we have not yet said what counts as "the" morphism v -> f. If there are multiple edges on the face that touch the vertex, there will be multiple composable paths from v to f. Do those correspond to distinct morphisms or to the same one? The answer is not forced by the above; it depends on whether we take the free category on the incidence graph or quotient it to thinness.

3.2 The free category on the incidence graph vs quotienting to thinness

Start with the directed incidence graph G(P):

- vertices of G(P) are the cells V(P) ⊔ E(P) ⊔ F(P),
- directed edges of G(P) are the generating incidences v -> e and e -> f.

The most literal categorical construction is the free category on this graph, call it Free(G(P)). In Free(G(P)):

- objects are the same as in G(P),
- morphisms are directed paths in G(P) (including length-0 paths as identities),
- composition is concatenation of paths.

This version remembers witness structure. A morphism is not merely the fact "v is incident to f," but the specific path through which that fact is witnessed. In particular:

- if a vertex v lies on a face f via two different boundary edges e1 and e2, then there are two distinct morphisms v -> f, namely the two paths v -> e1 -> f and v -> e2 -> f.

So Free(G(P)) is sensitive to multiplicity. It is not thin, and it contains more information than is typically used for duality as a poset reversal.

If our aim is duality as an "op" phenomenon, we often want to forget these multiplicities and keep only the yes/no incidence relation. This leads to the thin quotient.

Define a relation <= on cells by saying:

- a <= b iff there exists a morphism a -> b in Free(G(P)).

Because Free(G(P)) is built from incidence steps that only go upward in dimension, <= is a reachability relation. It is reflexive (there is always the identity path) and transitive (paths concatenate). Under the usual cell structure of a polyhedron, it is also antisymmetric when interpreted as face containment, yielding a partial order.

Now define the thin incidence category Thin(P) as the category associated to this preorder/poset:

- Ob(Thin(P)) = V(P) ⊔ E(P) ⊔ F(P),
- there is a unique morphism a -> b iff a <= b.

Equivalently: Thin(P) is obtained from Free(G(P)) by identifying all parallel morphisms between the same source and target. All distinct witnessing paths collapse into a single arrow representing mere reachability.

This thin quotient is exactly what the figure is morally displaying: it cares that a connection exists, not how many distinct paths witness it. The diagram shows many lines because there are many generating incidences, but it does not distinguish multiple composites with the same endpoints. In that sense, the figure is a rendering of the thin viewpoint even if it is drawn from generating incidence steps.

A useful way to summarize the relationship is:

- Free(G(P)) is the “path-sensitive” incidence category,
- Thin(P) is the “reachability-only” incidence category,
- there is a canonical functor q : Free(G(P)) -> Thin(P) that is identity on objects and sends every morphism/path to the unique arrow between the same endpoints.

This q is a quotient in the intuitive sense: it forgets witness structure.

For the remainder of the duality story, Thin(P) is typically the right object, because duality is about reversing the containment/incident order, and that order is inherently thin.

3.3 Equivalence to the face poset viewpoint

Classically, polyhedra are encoded by their face posets. The face poset Fpos(P) is the set of all faces of all dimensions ordered by inclusion. There are two common conventions:

- include the empty face and the whole polyhedron as top/bottom elements, or
- restrict to proper faces (vertices/edges/facets) depending on the application.

For the figure we are using the “3-stratum” version: vertices, edges, and faces only, no empty face and no whole polyhedron. But the identification with the face poset is straightforward.

Define Fpos_3(P) to be the poset whose underlying set is:

- V(P) ⊔ E(P) ⊔ F(P),

with order relation:

- a <= b iff a is a face of b (equivalently, a is contained in the closure of b).

Then the associated category Cat(Fpos_3(P)) is thin, with a unique morphism a -> b exactly when a <= b.

The claim is that Thin(P) as defined via reachability in the incidence graph coincides with Cat(Fpos_3(P)), under the standard assumption that incidence steps generate face containment. Intuitively:

- a vertex is contained in an edge exactly when it is incident as an endpoint,
- an edge is contained in a face exactly when it lies in its boundary,
- and containment between non-adjacent dimensions (vertex in face) is generated by chaining the adjacent containments.

Thus the thin incidence category is simply the face poset in categorical clothing.

This equivalence is not only conceptual; it is what makes the duality statement clean. Polyhedral duality is traditionally formulated as an order-reversing bijection between faces:

- phi : Fpos(P) -> Fpos(P*),
- with a <= b in P iff phi(b) <= phi(a) in P*.

That is exactly the assertion that:

- Fpos(P*) ≃ Fpos(P)^{op}

as posets, and therefore:

- Cat(Fpos(P*)) ≃ Cat(Fpos(P))^{op}

as thin categories.

So when we say “duality is contravariant,” the most precise meaning in this setting is: the incidence/face category of the dual is equivalent to the opposite of the incidence/face category of the original.

The figure uses the cube–octahedron pair to make this abstract statement tangible. It selects and draws one component of the order reversal explicitly: faces of the cube correspond to vertices of the octahedron. But the proper categorical claim is broader: there is a dimension-reversing correspondence on all faces that reverses incidence, and that reversal is exactly what the “op” construction formalizes.

In short, I(P) can be defined in two closely related ways:

- as a free category on the incidence graph (path-sensitive),
- or as the thin category associated to the face poset (order-sensitive).

The paper will primarily use the thin/poset version as the cleanest statement of duality, while keeping the free-category view in the background as an explanation of why “composition = derived incidence” is the natural mechanism behind the scenes.

4. The Cube C: Explicit Presentation of I(C)

To make the incidence category I(C) concrete, we choose a standard combinatorial model of the cube that makes every cell and every incidence relation explicit. The goal is not to “compute geometry,” but to exhibit a finite presentation of the cube’s incidence structure that can be read as a category: objects are cells, and arrows are incidence. The figure uses coordinate labels because they are maximally rigid: they prevent ambiguity about which vertex or face is which, and they make the cube’s symmetry visible without any additional explanation. Once this model is fixed, the incidence category is determined.

4.1 Vertex model V(C) = {±1}^3

We realize the cube’s vertex set as the eight sign triples:

V(C) = { (x,y,z) : x,y,z in {+1,-1} }.

Thus each vertex is a unique triple of signs, and there are exactly 2^3 = 8 vertices. This model is purely combinatorial: we may imagine these points in R^3, but the only structure we actually use is that each coordinate can flip independently.

The labeling in the left panel of the figure is exactly this: every blue dot is a vertex (±1,±1,±1). This gives two immediate benefits for I(C).

First, it provides a canonical description of the cube’s symmetries. Any permutation of coordinates and any independent sign flip acts on this set and corresponds to a symmetry of the cube. This matters later when we discuss why the duality map looks “coordinate aligned.”

Second, it makes “adjacency along an edge” entirely precise: vertices are adjacent if you can change exactly one sign and leave the other two unchanged. That adjacency rule is the next piece of the incidence category.

4.2 Edge adjacency as Hamming distance one

An edge of the cube connects two vertices that differ in exactly one coordinate. In the sign-triple model, define the Hamming distance between two vertices $u=(x1,y1,z1)$ and $v=(x2,y2,z2)$ as the number of coordinates in which they differ:

$d_H(u,v) = \#\{ i \text{ in } \{x,y,z\} : u_i \neq v_i \}$.

Then the cube’s edges are exactly the unordered pairs $\{u,v\}$ with

$d_H(u,v) = 1$.

Equivalently, we may write:

$E(C) = \{ \{u,v\} : u,v \text{ in } V(C), d_H(u,v)=1 \}$.

Counting check: from any vertex there are three neighbors (flip x, flip y, or flip z), so the total number of directed adjacency moves is $8*3 = 24$. Each undirected edge is counted twice, so $|E(C)| = 24/2 = 12$, as expected.

Categorically, edges become objects in I(C), but their incidence with vertices becomes generating morphisms. For each edge $e = \{u,v\}$, there are two vertex-edge incidence morphisms:

$u \to e$ and $v \to e$.

In the figure these are the solid segments from the blue vertex nodes in the leftmost column to the orange edge nodes in the middle column. The edge labels $E_0, \ldots, E_{11}$ are arbitrary names for these 12 objects. What matters is that the pattern of solid lines encodes the endpoint relation.

This is the first half of the generating incidence data:

- objects: $V(C) \sqcup E(C) \sqcup F(C)$,
- generating arrows: $v \to e$ if v is an endpoint of e.

To complete I(C), we need the faces and the edge-face incidence.

4.3 Face labels $F_{\pm x}$, $F_{\pm y}$, $F_{\pm z}$ and boundary incidence

In the coordinate model, the cube has six faces, each defined by fixing one coordinate to either +1 or -1. The figure labels them:

- F_{+x}, F_{-x},
- F_{+y}, F_{-y},
- F_{+z}, F_{-z}.

Formally, each face can be specified either as a set of vertices or as a set of edges (its boundary cycle). The vertex-set description is:

$F_{+x} = \{ (x,y,z) \text{ in } V(C) : x = +1 \}$,
$F_{-x} = \{ (x,y,z) \text{ in } V(C) : x = -1 \}$,

and similarly for y and z. Each of these sets has four vertices because fixing one coordinate leaves two free signs:

$|F_{\pm x}| = |\{(\pm 1,y,z) : y,z \text{ in } \{\pm 1\}\}| = 4$.

To define the edge-face incidence relation E -> F, we can state it in vertex terms:

An edge $e = \{u,v\}$ lies on a face F_{+x} if and only if both endpoints lie on that face, i.e.

$e \to F_{+x} \leftrightarrow (u \text{ in } F_{+x} \text{ and } v \text{ in } F_{+x})$.

In the sign-triple model, that condition has a simple meaning. Suppose e connects two vertices differing in exactly one coordinate. Then:

- If the differing coordinate is x, the edge crosses from $x=+1$ to $x=-1$, so it lies on neither F_{+x} nor F_{-x} but lies on two faces of type ±y and ±z determined by the fixed signs.
- If the differing coordinate is y, it lies on $F_{\pm x}$ and $F_{\pm z}$ determined by fixed signs.
- If the differing coordinate is z, it lies on $F_{\pm x}$ and $F_{\pm y}$ determined by fixed signs.

From this, a universal cube fact follows and is visible in the diagram: every cube edge lies on exactly two faces. Categorically, that means each edge object e has exactly two generating morphisms of type e -> f. In the figure these are the dashed segments from the edge nodes in the middle column to the face nodes in the right column.

At this point the incidence category I(C) is fully specified by generators:

- objects: the 8 vertices, 12 edges, 6 faces,
- generating morphisms: v -> e for endpoint incidence, e -> f for boundary incidence.

From these generators, all derived incidence morphisms appear by composition. For example, if v is an endpoint of e and e is in the boundary of face f, then there is a composite:

$v \to e \to f$,

which we interpret as a morphism $v \to f$ in the thin/poset version. In geometric language: "the vertex lies on the face." The cube makes the multiplicity issue especially vivid: each vertex lies on three faces, and on each such face there are two incident edges. In the free-path viewpoint this produces multiple paths from a given vertex to a given face. In the thin viewpoint these collapse to a single arrow whenever incidence holds.

So I(C) can be summarized as the categorical closure of two incidence relations:

- endpoint incidence $V(C) \to E(C)$,
- boundary incidence $E(C) \to F(C)$,

with the coordinate labeling giving a rigid, inspectable presentation.

This explicit cube presentation is not merely a warm-up. It sets up the later duality statement. The cube's coordinate faces $F_{\pm x},F_{\pm y},F_{\pm z}$ are exactly the objects that are sent, under duality, to the octahedron's coordinate vertices $V_{\pm x},V_{\pm y},V_{\pm z}$. Because we have fixed the cube model so concretely, the duality map in the figure is no longer a slogan. It is a function between explicitly named objects in two explicitly defined categories.

5. The Octahedron O: Explicit Presentation of I(O)

To complement the cube model, we present the octahedron O in a way that makes its incidence data as explicit and rigid as possible. The octahedron is the cube's dual, and the particular labeling used in the figure is not arbitrary: it is chosen so that the duality map $d : F(C) \leftrightarrow V(O)$ becomes literally "coordinate face goes to coordinate vertex." This section defines the cell sets V(O), E(O), F(O) and the incidence relations that generate I(O). The payoff is that the right panel of the figure becomes a faithful rendering of a concrete categorical object, rather than a schematic suggestion.

5.1 Vertex model $V(O) = \{\pm e_x, \pm e_y, \pm e_z\}$

The octahedron has six vertices. In the coordinate model, we take these vertices to be the six signed coordinate axes in R^3:

$V(O) = \{ (+1,0,0), (-1,0,0), (0,+1,0), (0,-1,0), (0,0,+1), (0,0,-1) \}$.

Equivalently, writing $e_x=(1,0,0)$, $e_y=(0,1,0)$, $e_z=(0,0,1)$, we have

$V(O) = \{ \pm e_x, \pm e_y, \pm e_z \}$.

The figure labels these vertices as

$V_{+x}, V_{-x}, V_{+y}, V_{-y}, V_{+z}, V_{-z}$, which you can interpret as shorthand for the six vectors above.

This model is dual-friendly in an immediate sense: the cube's six coordinate faces are "normal to" these six axis directions. The dotted correspondences in the center of the figure are exactly the identification of cube coordinate faces with octa coordinate vertices.

Once the vertex set is fixed, the edges and faces of the octahedron are determined by the combinatorics of which vertices can be connected without crossing the interior. The octahedron has 12 edges and 8 faces, and both numbers will emerge naturally from the sign-based description of faces.

5.2 Faces as sign-triples and how 8 arises naturally

A key fact about the octahedron is that each face is a triangle formed by choosing one vertex from each of the three axis pairs $\{\pm e_x\}$, $\{\pm e_y\}$, $\{\pm e_z\}$. In other words, to specify a face you choose a sign for x, a sign for y, and a sign for z. This yields a sign triple:

(s_x, s_y, s_z) in $\{+,-\}^3$.

Given such a triple, define the face (as a set of vertices) by:

$F_{s_x x + s_y y + s_z z} = \{ V_{s_x x}, V_{s_y y}, V_{s_z z} \}$.

For example:

- $F_{+x +y +z} = \{ V_{+x}, V_{+y}, V_{+z} \}$,
- $F_{+x +y -z} = \{ V_{+x}, V_{+y}, V_{-z} \}$,
- $F_{-x -y +z} = \{ V_{-x}, V_{-y}, V_{+z} \}$, etc.

There are $2^3 = 8$ sign triples, hence exactly **8 faces**. This is the conceptual mirror of the cube's $2^3 = 8$ vertices. In the figure, these eight faces appear in the right panel's face column as labels like F_{+x+y+z}, F_{+x+y-z}, and so on.

This sign-triple description also has a clean geometric interpretation that we can state without leaning on picture intuition. Each face lies in a plane of the form:

$s_x X + s_y Y + s_z Z = 1$,

where (X,Y,Z) are coordinates in R^3, and the three chosen vertices indeed satisfy that equation. For instance, the three vertices (1,0,0), (0,1,0), (0,0,1) all satisfy X+Y+Z=1, so they form the face $F_{+x +y +z}$. Similarly, (1,0,0), (0,1,0), (0,0,-1) satisfy X+Y-Z=1, yielding $F_{+x +y -z}$. The entire octahedron can be described as:

$O = \{ (X,Y,Z) \text{ in } R^3 : |X| + |Y| + |Z| <= 1 \}$, and its faces are precisely the planar patches where the absolute-value expression becomes linear with a fixed sign pattern.

From the categorical viewpoint, what matters is that the face set F(O) is canonically indexed by {+,-}^3, so the number 8 is not a memorized fact; it is an inevitable consequence of choosing one sign for each axis.

Now we define edges by incidence between these vertices and faces.

An octahedron edge connects two vertices that are not opposite. In label language:

V_a is opposite V_b if b = -a (e.g., +x opposite -x).

Then edges are exactly pairs of distinct non-opposite vertices:

E(O) = { {V_a, V_b} : a != b and a != -b }.

Counting check: each vertex is connected to four others (everything except itself and its opposite). Thus directed adjacency count is 6*4 = 24, divide by 2 gives |E(O)| = 12.

So, in the incidence category I(O):

- the generating morphisms v -> e correspond to endpoint incidence,
- the generating morphisms e -> f correspond to boundary incidence.

Using the face-as-vertex-triple description, edge-face incidence is easy to state:

An edge {V_a,V_b} lies on a face F_{s_x x + s_y y + s_z z} iff both endpoints are among that face's three vertices. Equivalently:

{V_a,V_b} -> F_{s_x x + s_y y + s_z z} <-> V_a in F_{...} and V_b in F_{...}.

Because each face has exactly three vertices, each face has exactly three edges, and the categorical pattern "each face has three incoming dashed edges from E" follows immediately.

5.3 Matching incidence patterns to the diagram

With the explicit presentation in hand, we can now interpret every visible structural feature of the right panel as a categorical fact about I(O).

First, the object counts match the known combinatorics:

- |V(O)| = 6 (six blue dots in the vertex column),
- |E(O)| = 12 (twelve orange squares in the edge column),
- |F(O)| = 8 (eight green triangles in the face column).

Second, the degrees (how many morphisms touch each object) match the octahedron's local structure.

Vertex-to-edge incidence:

- Each vertex is adjacent to four edges. Categorically, each $V_{\pm x}$, $V_{\pm y}$, $V_{\pm z}$ has exactly four outgoing generating morphisms V -> E.
- This is visible as four solid lines leaving each blue dot in the right panel.

Edge-to-face incidence:

- Each edge lies on exactly two faces. Categorically, each edge object has exactly two outgoing generating morphisms E -> F.
- This is visible as each orange square sending exactly two dashed lines to the face column.

Face incidence:

- Each face is a triangle with three edges and three vertices. Categorically, each face object has exactly three incoming E -> F morphisms, and via composition it has exactly three incoming V -> F morphisms in the thin version.
- In the diagram, each face triangle receives exactly three dashed lines (one from each of its boundary edges).

Third, the labeling scheme aligns with the intended duality. The faces labeled by sign triples $F_{s_x x + s_y y + s_z z}$ are precisely the objects that are naturally indexed by $V(C) = \{\pm 1\}^3$. This is why later, when we discuss the full duality correspondence, we can say without handwaving:

- cube vertices (sign triples) correspond to octa faces (sign triples).

The figure only draws the complementary correspondence:

- cube faces (six coordinate faces) correspond to octa vertices (six axis vertices).

But the right panel's face labels are already preparing the ground for the omitted eight dotted lines that would connect V(C) to F(O) if we chose to draw them. The diagram is therefore internally consistent: the same sign-triple language appears on the cube's vertex column and the octahedron's face column, and the same axis-pair language appears on the cube's face column and the octahedron's vertex column.

Finally, the categorical closure under composition is implicitly present on the right panel exactly as it was on the left. Whenever the diagram exhibits V -> E and E -> F, the category contains composites V -> F. In geometric terms, "vertex lies on face" is derived from "vertex lies on edge" and “edge lies on face.” In the free-path viewpoint, there are typically multiple distinct paths witnessing a given vertex-face relation. In the thin/poset viewpoint, there is at most one arrow V -> F for a given pair of objects, and it exists precisely when the vertex is a member of the face’s vertex triple.

So the right panel is not merely “an octahedron drawn categorically.” It is an explicit, checkable presentation of I(O) in which:

- vertices are axis basis points,
- faces are sign triples,
- edges are non-opposite pairs,
- and incidence becomes the generating morphisms that are then closed under composition.

This completes the symmetry with the cube and sets the stage for the next section, where duality becomes a contravariant statement about how I(C) and I(O) relate as categories (and in the thin case, as opposite categories).

6. Duality as Contravariance: The Opposite-Category Principle

The heart of the paper is a single structural claim: polyhedral duality is naturally contravariant. In ordinary geometry this is often stated informally as “faces become vertices and vertices become faces.” Category theory sharpens that into a statement about reversed arrows and opposite categories. The “opposite-category principle” is the idea that the correct categorical home for duality is not a map I(P) -> I(P*) that preserves arrow direction, but a map I(P) -> I(P*)^{op} (or equivalently I(P*) -> I(P)^{op}) because incidence reversal is built into the duality operation itself. The figure embodies this: the incidence arrows inside each panel point upward in dimension, while the dotted duality correspondences are meant to be read as a dimension-reversing identification.

6.1 Why duality reverses incidence directions in general

Duality is dimension-reversing, but more specifically it is incidence-reversing. The cleanest way to say this is:

- in a polyhedron, a lower-dimensional cell is “contained in” or “incident to” a higher-dimensional cell,
- in the dual polyhedron, the corresponding cells have swapped dimension, so the direction of containment is reversed.

A standard geometric construction makes this intuitive. For a convex polyhedron P, each face f determines a supporting plane. In the dual, a vertex v_f is placed corresponding to that supporting plane (often represented by the outward normal or by the plane's equation in a suitable polar dual). Then:

- if a vertex v lies on a face f of P, that means v satisfies the plane constraint defining f,
- in the dual, the dual face corresponding to v is the collection of dual vertices corresponding to faces that contain v.

Thus containment/incidence in P becomes reverse containment/incidence in P*. The reversal is not optional; it is the content of the construction. A face in P becomes a vertex in P*, and the set of faces containing a given vertex in P becomes the set of vertices on a corresponding face in P*.

At the level of "arrows," the simplest instance is:

- in P we have an incidence v -> e (vertex incident to edge),
- under duality, the duals are v* (a face) and e* (an edge),
- and the corresponding relation is e* -> v* (edge incident to face), i.e. reversed.

So if we view incidence as directed upward in dimension (as we did in the figure), then duality flips directions: it takes an incidence arrow and turns it into an arrow that points the other way when expressed among dual cells.

This is exactly what category theory calls contravariance: a structure-preserving map that reverses arrows is naturally a functor into an opposite category.

*6.2 Face posets and the relationship F(P) ≃ F(P)^{op}**

The most standard formal setting for polyhedral duality is the face poset. Let F(P) denote the set of faces of P of all dimensions (often including the empty face and the whole polyhedron if one wants a bounded lattice). The poset order is inclusion:

- a <= b iff face a is contained in face b.

Duality provides a bijection between faces of P and faces of P* that reverses inclusion. Concretely, if a is a k-face of P, then a* is an (n-1-k)-face of P* (here n=3 for polyhedra, but the statement is dimension-general). The defining property is:

- a <= b in F(P) iff b* <= a* in F(P*).

This is order reversal. In poset language, that means:

- $F(P^*)$ is isomorphic to $F(P)^{op}$,

where $F(P)^{op}$ is the same underlying set of faces but with the order reversed.

Categorically, any poset can be viewed as a thin category: there is a unique morphism $a \to b$ iff $a \leq b$. Under this identification, order reversal becomes an equivalence of categories:

- $\mathrm{Cat}(F(P^*)) \simeq \mathrm{Cat}(F(P)^{op}) = \mathrm{Cat}(F(P))^{op}$.

So the "opposite-category principle" is not a philosophical add-on; it is exactly the categorical translation of the classical duality theorem for face posets.

In 3D, this gives the familiar swaps:

- vertices of P correspond to faces of P*,
- edges correspond to edges,
- faces correspond to vertices,

and all incidence is reversed because all inclusion is reversed.

It is worth emphasizing what this does and does not claim. It does not claim that $F(P) \simeq F(P)^{op}$ for a single polyhedron. That would mean P is self-dual at the level of face posets. Some polyhedra have this property (the tetrahedron is the classical example), but in general it is false. The correct general statement is:

- $F(P^*) \simeq F(P)^{op}$.

The star belongs on P, not inside the op. Duality relates P to its dual, and the relation is order-reversing.

6.3 Translating "op" into the three-stratum diagram

The figure uses a three-stratum layout, separating the objects by dimension into columns:

- V (0-cells),
- E (1-cells),
- F (2-cells),

and drawing generating incidence morphisms upward:

- $V \to E$,
- $E \to F$.

This layout is already a categorical choice: it treats incidence as directed and composable.

Now, what does op mean in this picture language?

Taking the opposite category I(P)^{op} means:

- keep the same objects,
- reverse every arrow.

So in the diagrammatic vocabulary:

- every solid incidence line v -> e becomes e -> v,
- every dashed incidence line e -> f becomes f -> e,
- and therefore every composite v -> f becomes f -> v.

If we tried to draw I(P)^{op} using the same left-to-right column order V, E, F, all arrows would point downward (from higher dimension to lower). That is one legitimate convention, but the figure chooses another strategy: it keeps each panel internally "upright" (arrows go upward), and expresses duality not by flipping arrows inside a panel, but by pairing I(C) with I(O) so that "upward in one corresponds to upward in the other after a dimension swap."

Concretely, the duality correspondence has the effect:

- F(C) corresponds to V(O),
- E(C) corresponds to E(O),
- V(C) corresponds to F(O).

If we write this as an object mapping D (suppressing technicalities about whether it lands in I(O) or I(O)^{op}), then the key consistency condition is:

- an incidence arrow in I(C) becomes an incidence arrow in the opposite direction among dual objects.

For example:

- in I(C) there is an arrow e -> f whenever edge e lies on face f,
- under duality, f corresponds to a vertex D(f) in O, and e corresponds to an edge D(e) in O,
- the reversed incidence is D(f) -> D(e) in I(O) (vertex incident to edge).

So “edge-to-face” in the cube becomes “vertex-to-edge” in the octahedron, which is exactly a reversal of direction once you identify faces with vertices.

This is why the three-stratum diagram is such a good stage for contravariance. The dimension swap is visible. The arrow direction is visible. And the opposite-category principle becomes an almost literal rule for translating one panel’s incidences into the other panel’s incidences.

The dotted lines in the center are a partial manifestation of that translation. They show the object-level component:

- d : F(C) <-> V(O).

But the categorical content is richer: if we drew all object correspondences (including V(C) <-> F(O) and E(C) <-> E(O)), we could then check arrow reversal by inspection. Every solid line on one side would correspond to a dashed line on the other side after the appropriate object identification, and vice versa, reflecting exactly the fact that the functor implementing duality is contravariant.

So the three-stratum diagram translates op into a visual rule:

- duality swaps the V and F strata,
- keeps the E stratum matched,
- and reverses incidence direction, which is naturally expressed by landing in an opposite category.

This is the conceptual hinge of the entire paper. Once you accept that duality is best understood as a map into an opposite category, the cube–octahedron relationship stops being “a coincidence of numbers” and becomes a structured equivalence of incidence data. The remaining sections will make this precise in the cube–octahedron case (by specifying the full correspondence and verifying incidence preservation) and then explain how the same reasoning scales to general convex polyhedra and higher-dimensional polytopes.

7. The Duality Map d in the Figure: Faces of C to Vertices of O

The figure makes one portion of cube–octahedron duality explicit: a map from the six cube faces to the six octahedron vertices. This is the simplest piece of the duality because it lines up perfectly with the coordinate descriptions we chose. A cube face is a coordinate plane slice ($x=\pm1$, $y=\pm1$, $z=\pm1$), and the octahedron’s vertices are the six signed coordinate axes ($\pm e_x$, $\pm e_y$, $\pm e_z$). The map drawn by the six dotted lines is therefore the “coordinate duality” dictionary. But the categorical content of that dictionary is not merely a list of correspondences. It is a witness that incidence is being reversed in the expected way, and it is a fragment of a larger contravariant equivalence that also includes cube vertices and cube edges.

7.1 The concrete correspondence $d(F_{\pm x}) = V_{\pm x}$, etc.

In the coordinate cube model, the cube faces are labeled:

$F(C) = \{ F_{+x}, F_{-x}, F_{+y}, F_{-y}, F_{+z}, F_{-z} \}$.

In the coordinate octahedron model, the octahedron vertices are labeled:

$V(O) = \{ V_{+x}, V_{-x}, V_{+y}, V_{-y}, V_{+z}, V_{-z} \}$.

The map drawn in the figure is the bijection:

$d : F(C) \to V(O)$

defined by the six assignments:

- $d(F_{+x}) = V_{+x}$,
- $d(F_{-x}) = V_{-x}$,
- $d(F_{+y}) = V_{+y}$,
- $d(F_{-y}) = V_{-y}$,
- $d(F_{+z}) = V_{+z}$,
- $d(F_{-z}) = V_{-z}$.

In vector terms, if we identify V_{+x} with $(+1,0,0)$ and so on, then:

$d(F_{\pm x}) = \pm e_x$, $d(F_{\pm y}) = \pm e_y$, $d(F_{\pm z}) = \pm e_z$.

One can motivate this assignment geometrically (face normals, supporting planes, polarity), but for the purpose of the figure it is enough that it is the unique “coordinate-consistent” identification: the cube face defined by fixing x is paired with the octa vertex lying on the x axis.

Categorically, this map is at the object level: it does not yet say anything about morphisms. Its role is to align strata across the two incidence categories:

- cube faces (2-cells) are being identified with octa vertices (0-cells).

That is precisely one component of the general dimension reversal rule for duality in 3D.

7.2 What the six dotted lines certify

The six dotted lines are not incidence arrows inside either panel. They are cross-panel correspondences. Their job is to certify three things simultaneously.

First, they certify that the figure is not merely “showing two categories side by side.” It is showing a relationship between them. Without the dotted lines, you could interpret the figure as two unrelated incidence categories. With them, the viewer is forced to read the two panels as a dual pair, tied together by an explicit mapping.

Second, the dotted lines certify the correct count-level symmetry of duality at this stratum: $|F(C)| = 6$ and $|V(O)| = 6$. This count equality is not the main theorem, but it is a necessary sanity check. The dotted lines show a bijection, not just a loose association.

Third, and most importantly, they certify the intended direction of contravariance when you compare incidence patterns on the two sides. Even though the dotted lines themselves are not arrows, they enable a “transport of structure” test: you can use them to translate a cube face into an octa vertex and then ask whether the surrounding incidences reverse as they should.

Here is the essential incidence-reversal check that these correspondences support.

On the cube side, take a cube face, say F_{+x}. Consider the set of cube edges incident to that face. In I(C), those incidences appear as dashed arrows:

e -> F_{+x}

for each boundary edge e of that face.

Under duality, F_{+x} is mapped to the octa vertex V_{+x}. In the octahedron incidence category I(O), the edges incident to V_{+x} appear as solid arrows:

V_{+x} -> e'

for each octa edge e' adjacent to that vertex.

Thus, the “edges incident to the face” on the cube side should correspond to the “edges incident to the vertex” on the octa side, but with direction reversed. The dotted line enables you to line up the face node on the cube panel with the vertex node on the octa panel and visually compare their local incidence neighborhoods.

The figure, as drawn, does not explicitly label the edge correspondence E(C) <-> E(O), so you cannot complete the full transport check solely by inspection. But you can already see the pattern: cube faces have four boundary edges; octa vertices have four incident edges. The valence matches exactly. That match is not accidental: it is one of the cleanest numerical signatures of duality.

So the dotted lines certify that the diagram is arranged consistently with contravariance:

- a “face-adjacent dashed neighborhood” in the cube is meant to correspond to a “vertex-adjacent solid neighborhood” in the octahedron.

This is the diagrammatic translation of “incidence direction reverses under duality.”

7.3 The missing correspondences: V(C) ↔ F(O) and E(C) ↔ E(O)

The full cube–octahedron duality is not just the face-to-vertex map. In 3D it has three components:

- F(C) <-> V(O) (drawn),
- E(C) <-> E(O) (not drawn),
- V(C) <-> F(O) (not drawn).

The figure omits the latter two correspondences for a practical reason: drawing all of them would crowd the center and reduce legibility. But mathematically, these omitted correspondences are the majority of the structure. They are also the ones that make the contravariant equivalence statement complete.

The vertex–face correspondence V(C) <-> F(O) is particularly elegant in our chosen labeling schemes. Cube vertices are sign triples:

$v = (s_x, s_y, s_z)$ with each s_i in $\{+1,-1\}$.

Octa faces are also indexed by sign triples:

$F_{s_x x + s_y y + s_z z}$ with each s_i in $\{+,-\}$.

So we can define a map:

$d_V : V(C) \to F(O)$

by:

$d_V((s_x, s_y, s_z)) = F_{(\mathrm{sign}(s_x))\, x + (\mathrm{sign}(s_y))\, y + (\mathrm{sign}(s_z))\, z}$.

In words: the cube vertex with sign pattern (±,±,±) maps to the octa face whose three vertices are the corresponding axis vertices $V_{\pm x}$, $V_{\pm y}$, $V_{\pm z}$.

This is not an arbitrary bijection. It is the canonical one induced by the coordinate models: a cube vertex is a choice of signs for the three coordinates, and an octa face is also a choice of signs for the three coordinate axes. Hence both sets have size 8 and carry the same $\{\pm\}^3$ indexing, making the correspondence immediate.

The edge correspondence E(C) <-> E(O) is slightly subtler to describe, but it is equally canonical. Under duality in 3D, edges map to edges. In the coordinate models, you can describe a cube edge by:

- fixing two coordinates and allowing the remaining coordinate to vary between -1 and +1.

For example, the edge with y=+1 and z=-1 and x varying connects (-1,+1,-1) to (+1,+1,-1).

On the octahedron side, an edge is a pair of non-opposite axis vertices, such as {V_{+x}, V_{+y}} or {V_{+x}, V_{-z}}.

The conceptual rule for matching them is “boundary adjacency becomes vertex adjacency.” An edge in the cube is the intersection of exactly two faces, and in the dual octahedron it should correspond to the edge that connects the two dual vertices of those faces. Formally:

If a cube edge e lies on exactly two faces f1 and f2, then its dual edge is:

d_E(e) = { d(f1), d(f2) },

an edge connecting the two octa vertices corresponding to those two cube faces.

This is the cleanest categorical description because it uses only incidence information. It does not require geometry. It says: edges are “codimension-1 intersections of faces,” and duality turns “faces” into “vertices,” so “intersection of two faces” becomes “edge connecting two vertices.”

Once d_E is defined this way, the contravariance check becomes a real statement:

- in I(C), vertex-edge incidence v -> e should correspond (under duality) to edge-face incidence d_E(e) -> d_V(v) in I(O),
- in I(C), edge-face incidence e -> f should correspond to vertex-edge incidence d(f) -> d_E(e) in I(O),

which is precisely “arrows reverse,” expressed after swapping strata.

So the full structure that the figure is compressing is:

- an object-level bijection on each stratum (with dimension reversal V <-> F and E <-> E),
- and an incidence reversal rule that can be stated as an isomorphism/equivalence between I(C) and I(O)^{op} in the thin/poset setting.

The six dotted lines are therefore best read as an intentionally minimal “handle” on the full duality: they pin down the coordinate alignment and allow the reader to infer how the rest of the correspondence must behave. The next section will make that inference explicit by stating what must be checked to upgrade these correspondences from a suggestive diagram into a verified contravariant equivalence of incidence categories.

8. Functorial Integrity: What Must Be Checked for a True Equivalence

The figure is persuasive because it compresses a lot of correct structure into a single page. But category theory forces a discipline that pictures alone cannot supply: if we claim that duality is an equivalence of incidence categories (up to opposite), then we must specify a functor (or contravariant functor) and verify that it actually preserves the categorical structure it is supposed to preserve. This section is the "audit layer." It separates what is easy to assert at the level of objects from what must be verified at the level of morphisms, and it spells out the minimal commuting conditions that convert a visually plausible correspondence into a true categorical equivalence.

8.1 Object-level bijections vs morphism-level preservation

The easiest part of any would-be equivalence is to match objects. In the cube–octahedron case, the object-level correspondences are familiar and can be stated as bijections of underlying sets:

- $d_F : F(C) \to V(O)$ (drawn in the figure),
- $d_V : V(C) \to F(O)$ (suggested by the sign-triple labels),
- $d_E : E(C) \to E(O)$ (determined by "edge is intersection of two faces").

At the object level, these bijections can be correct while the categorical claim is still false. The reason is that category structure is not "just objects." It is objects plus morphisms plus composition. A functor is not a bijection on objects. A functor must map morphisms to morphisms in a way that respects identities and composition.

In the thin/poset version, there is at most one morphism between any two objects. That makes life easier, but it does not eliminate the need for verification. In a thin category, a functor is essentially the same thing as an order-preserving map: it must send comparable pairs to comparable pairs. A bijection on objects that fails to preserve the order is not a functor, even if it "looks plausible."

In the cube–octahedron setting, we want contravariance: duality should reverse incidence. That means the clean categorical target is an opposite category. A precise statement looks like:

- there exists a functor $D : I(C) \to I(O)^{op}$

(or equivalently $D : I(C)^{op} \to I(O)$), such that D is an equivalence of categories.

At the object level, D must implement the dimension reversal:

- vertices of the cube go to faces of the octahedron,
- edges go to edges,
- faces go to vertices.

But an object map with these properties is not yet “duality” in the categorical sense. It becomes duality only if it also carries every generating incidence morphism to the corresponding generating incidence morphism (with direction reversed because of op), and does so coherently with composition.

So the audit question is:

- do our object-level bijections extend to a functor that maps incidence arrows to incidence arrows, respecting how incidence composes?

8.2 Incidence preservation as a commuting condition on composites

The figure shows two kinds of generating morphisms in each incidence category:

- V -> E (vertex incident to edge),
- E -> F (edge incident to face).

And it implicitly relies on the fact that composites exist:

- V -> E -> F gives a derived V -> F relation.

To say that a contravariant duality is functorial is to say, roughly:

- incidence chains in C correspond to reversed incidence chains in O.

This can be expressed as a commuting condition that is particularly clean in the thin/poset formulation.

Let <=_C denote the incidence/containment order in I(C) and <=_O the order in I(O). In thin categories, “there is a morphism a -> b” is equivalent to “a <= b.” Contravariance then means order reversal:

- for all cells a,b of C,
  a <=_C b iff D(b) <=_O D(a).

This single biconditional is the entire functoriality condition at once: it says D is an order anti-isomorphism between the face posets, hence a functor into the opposite category.

But the figure is drawn in three strata using only the adjacent-dimension incidences as visible generators. So it is useful to state the check in those terms as well.

The generating incidence checks are:

- if v -> e in I(C), then D(e) -> D(v) in I(O) (because arrows reverse),
- if e -> f in I(C), then D(f) -> D(e) in I(O).

Notice the swap:

- v -> e becomes edge -> face on the dual side (since D(v) is a face and D(e) is an edge),
- e -> f becomes vertex -> edge on the dual side (since D(f) is a vertex and D(e) is an edge).

Now the key “commuting composites” check is: whenever two arrows compose on the cube side, their images must compose (in reversed order) on the octa side and land at the same endpoints.

Take a cube incidence chain:

v -> e -> f.

Under a contravariant functor D : I(C) -> I(O)^{op}, this should map to a chain in I(O) of the form:

D(f) -> D(e) -> D(v).

This is the diagrammatic meaning of functoriality: the image of the composite equals the composite of the images, taking into account reversal. Written as a commuting condition:

- D(v -> e -> f) = D(f) -> D(e) -> D(v),

where the left side is the image of the composite morphism in I(C) and the right side is the composite of the images in the appropriate category. In a thin category this reduces to existence: if the left-side composite exists, then the right-side composite must exist, and conversely.

In practical terms, this means:

- If a cube vertex lies on a cube face (witnessed by some boundary edge), then the dual octa vertex must lie on the dual octa face.
- Conversely, if an octa vertex lies on an octa face, the corresponding cube face must contain the corresponding cube vertex.

This is where the sign-triple labeling becomes more than aesthetic. It makes these commuting conditions checkable by “symbol pushing” rather than by geometric intuition. For instance, the cube vertex (s_x,s_y,s_z) is contained in the cube face F_{+x} exactly when $s_x = +1$. Under the duality dictionary, this becomes a statement that the octa vertex V_{+x} lies on the octa face $F_{s_x x + s_y y + s_z z}$ exactly when the face's vertex set includes V_{+x}, which is again "exactly when $s_x=+1$." The condition matches by construction. That is functorial integrity in its simplest explicit form.

8.3 When the diagram witnesses an equivalence and when it doesn't

A diagram can be correct as a visualization and still fail to witness an equivalence if it does not encode enough structure, or if it encodes structure incorrectly. Here are the main success and failure modes.

The diagram witnesses a true equivalence (up to opposite) when:

- The object correspondences are bijections on each stratum, and they are consistent with dimension reversal (V <-> F, E <-> E, F <-> V).
- For every generating incidence arrow shown in the cube panel, there is a corresponding generating incidence arrow on the octa panel after translating endpoints via the object map and reversing direction.
- The same holds in reverse: every translated octa incidence corresponds back to a cube incidence.
- The derived incidences (composites) are preserved automatically because the generators are preserved and composition is forced.

In the thin/poset setting, these conditions collapse to: the correspondence is an order anti-isomorphism between face posets, hence it defines an isomorphism of thin categories to an opposite category.

The diagram does not witness an equivalence when:

- The object matching is only partial and the omitted parts matter for verifying arrow reversal. In our figure, only F(C) <-> V(O) is explicitly drawn. The other correspondences exist implicitly but are not shown. That means the figure strongly suggests equivalence, but on its own it does not *certify* it unless the reader supplies the missing matches from the labels.
- The object matching is a bijection but is "wrong" relative to incidence. For instance, you could permute the six cube faces arbitrarily and still get a bijection to the six octa vertices. Most such bijections will not preserve adjacency patterns: a cube face has four boundary edges, so its image vertex must have four incident edges; that degree

constraint is necessary but not sufficient. If you mismatch which faces meet along which edges, the functoriality conditions fail.

- The diagram uses a non-thin interpretation (path-sensitive category) but the intended duality only respects the thin quotient. In a free-path category, there can be multiple distinct morphisms v -> f corresponding to different boundary edges. A naive duality map that collapses those distinctions might fail to be fully faithful if you insist on the richer morphism sets. Put bluntly: if you keep witness paths, then you must decide what duality does to those witness paths. If you do not, you have not defined a functor at that level.
- The figure conflates “graph isomorphism” with “categorical equivalence.” Two incidence graphs can look similar while differing in how compositions are identified or in whether certain morphisms are forced to be equal. Category theory makes those distinctions explicit. An equivalence requires the right preservation of composition, not merely the right adjacency pattern.

For our purposes, the diagram is best understood as a faithful witness of equivalence in the thin/poset sense, provided we supply the object-level correspondences that the labels already encode. The dotted lines show one component of the object map; the sign-triple labels on cube vertices and octa faces give the second; and the “edge as intersection of two faces” rule gives the third. Once these are fixed, the commuting incidence conditions become checkable facts rather than aesthetic impressions.

This section therefore sets the standard for the remainder of the paper: whenever we claim “duality is an equivalence (to an opposite category),” we mean that there exists a precisely defined object mapping and an induced morphism mapping such that incidence arrows correspond under reversal, and composites correspond as required. The next section generalizes this audit logic beyond the cube–octahedron pair, showing that the same functorial integrity conditions can be formulated and verified for any convex polyhedron and its dual.

9. Generalization Blueprint: Any Convex Polyhedron and Its Dual

The cube–octahedron pair is the cleanest “page-sized” example because the labeling can be made coordinate-rigid, the duality map has an obvious description, and the incidence degrees line up in a way that is easy to see. But the categorical story is not special to Platonic solids. It is a general feature of convex polytopes: duality is an order-reversing correspondence between faces, and therefore an equivalence between face-incidence categories up to opposite. This section outlines a reusable blueprint for moving from the cube–octahedron picture to an arbitrary convex polyhedron P (and beyond), without relying on any special symmetry.

The guiding theme is: geometry provides a way to construct the dual, but the duality theorem lives in combinatorics. Once you translate a polyhedron into its face-incidence order, duality becomes a statement about posets (hence thin categories), and the "opposite-category principle" remains the correct structural form in all dimensions.

9.1 Face lattices, polytopes, and combinatorial duality

For a convex polytope P (in any dimension), consider its set of faces of all dimensions, including (if desired) the empty face and the polytope itself. Let F(P) be that set, ordered by inclusion:

- a <= b iff a is a face of b (equivalently, a is contained in b as a subset of R^n).

This is the face poset of P. Under mild assumptions, the face poset has additional structure: it is graded by dimension (or rank), and in many treatments, when one adjoins bottom/top elements, it becomes a lattice (the face lattice). The details of lattice operations are not necessary for our purpose; what matters is that the face structure is organized by inclusion, and inclusion is exactly what the incidence arrows encode.

Convex duality, in its standard polar form, produces a dual polytope P* such that faces of P correspond to faces of P* in a dimension-reversing way. The fundamental combinatorial theorem can be stated purely in poset language:

- there exists an order-reversing bijection phi : F(P) -> F(P*) such that
  a <= b in F(P) iff phi(b) <= phi(a) in F(P*).

This is precisely the statement:

- F(P*) ≃ F(P)^{op}

as posets.

Categorically, if we regard F(P) as a thin category (unique morphism a -> b iff a <= b), then the same statement becomes:

- Cat(F(P*)) ≃ Cat(F(P))^{op}.

This is the general version of what the cube–octahedron figure is expressing in three strata. The cube–octahedron diagram simply restricts attention to the faces of ranks 0,1,2 (vertices, edges, facets) and draws the generating cover relations among them. The blueprint for general P is therefore:

1. define the face poset F(P) (or its truncated version if you only want certain ranks),
2. interpret it as a thin category,
3. define the dual P* and the duality map on faces,

4. verify that the map is order-reversing (equivalently functorial into an opposite category).

Once that is done, any “category picture” is just a visualization of this poset-category structure.

A practical note: for many polyhedra, the easiest way to compute or represent the face poset is not geometric but combinatorial: list vertices, edges, facets and their containment relations. This is why the category-theoretic approach is amenable to computation: the input is finite incidence data, and the output is a small category (often thin) that can be drawn.

9.2 Extending to higher dimensions: k-faces and rank reversal

Everything above becomes more systematic in higher dimensions because the language of ranks (dimensions) generalizes cleanly.

Let P be a convex polytope in R^n. A face has a dimension k with -1 <= k <= n, where -1 is the empty face and n is P itself (if included). The face poset is graded by this dimension; equivalently, by rank:

- rank(face) = dim(face) + 1 if one wants the empty face to have rank 0.

Duality in n dimensions reverses ranks:

- a k-face of P corresponds to an (n-1-k)-face of P*.

In 3D, this is the familiar 0 <-> 2 swap and 1 <-> 1 fix. In 4D, it becomes:

- vertices (0-faces) <-> 3-faces (facets),
- edges (1-faces) <-> 2-faces (ridges),
- and so on.

Categorically, this is still contravariance: inclusion is reversed. The same statement holds:

- F(P*) ≃ F(P)^{op}.

But now, if one wants a “stratified diagram,” there will be n+1 strata of objects (for each face dimension), and generating morphisms typically correspond to cover relations in the graded poset:

- a -> b when a < b and dim(b) = dim(a)+1 and there is no face strictly between them.

In other words, the natural directed graph to draw is the Hasse diagram of the face poset, with edges oriented upward in rank. A “category picture” in higher dimensions is exactly a Hasse diagram of F(P) (or a slice of it) understood as a thin category.

Duality then becomes: reverse the diagram and relabel by dual faces. Graphically, duality corresponds to flipping the stratification: the lowest stratum becomes the highest, and cover relations reverse direction. Categorically, this is precisely “pass to the opposite category.”

This is one reason the opposite-category principle is so apt: it is not a special property of 3D polyhedra; it is the universal property of face inclusion under duality.

9.3 What changes for non-simple or non-simplicial polytopes

The cube–octahedron pair has an extra niceness: one is simple and the other is simplicial.

- A polytope is **simple** if exactly n edges meet at each vertex (in n dimensions).
- A polytope is **simplicial** if every facet is a simplex (triangular in 3D, tetrahedral in 4D, etc.).

In 3D:

- the cube is simple: 3 edges meet at every vertex,
- the octahedron is simplicial: every face is a triangle.

Duality swaps these properties:

- the dual of a simple polytope is simplicial, and vice versa.

This matters for how “uniform” the incidence degrees look in a diagram. In the cube panel, every vertex has degree 3 in the V -> E incidence. In the octa panel, every face has exactly 3 boundary edges in the E -> F incidence. That uniformity is a hallmark of simple/simplicial structure and contributes to the aesthetic clarity of the figure.

For a general polyhedron, especially one that is neither simple nor simplicial, the local incidence degrees vary:

- vertices can have different valences (numbers of incident edges),
- faces can have different numbers of sides (numbers of boundary edges).

Duality still exists and still reverses incidence, but the diagram will no longer be “regular.” Instead, the diagram will show that variability as an intrinsic part of the incidence category. Categorically, nothing breaks. A poset-category does not require uniform degrees; it only requires a well-defined order relation. Duality remains an anti-isomorphism of face posets. What changes is the visual symmetry and the ease of recognition.

There is also a structural subtlety: in non-simple or non-simplicial cases, there are more “short” incidence chains that can witness the same derived relation. For example, a vertex can lie on a face in many boundary edges; a face can meet another face along more complicated intersections in higher dimensions. In the thin category, all those witnesses collapse. In the free-path category, they remain distinct. This is exactly where the choice of category model matters:

- If your goal is duality as order reversal, the thin face-poset category is stable and canonical.
- If your goal is to retain witness multiplicity, you must decide what duality does to those witnesses, which can be more intricate in non-regular polytopes.

Finally, in practice, representing F(P) explicitly becomes harder as complexity grows. Even in 3D, a polyhedron with many faces can yield a large incidence category. In higher dimensions, the number of faces of all ranks can explode. This is where computational methods and selective visualization become essential: one may draw only certain ranks, or only the Hasse diagram restricted to a subposet, or only the adjacency graph of facets, depending on what one wants to communicate.

So the generalization blueprint is robust:

- define the face-incidence category (typically thin),
- define the duality map on faces,
- interpret duality as an equivalence to an opposite category.

What changes from the Platonic case is not the theorem but the "regularity of the picture." The cube–octahedron figure is an ideal demonstration because it is small, symmetric, and labelable. But the categorical mechanism it exhibits is universal: duality is contravariant because inclusion/incident structure is reversed, and the opposite category is the natural place where that reversal becomes a structure-preserving map.

10. Relation to HoTT and Higher Categories

Up to this point, the paper has deliberately stayed in the safe zone: thin categories (posets) built from incidence, and duality understood as an order-reversing equivalence, i.e. a functor into an opposite category. That story is already complete and already powerful. So why mention HoTT (Homotopy Type Theory) or higher categories at all?

Because the moment you stop collapsing "witnesses" and start caring about *how* incidence is witnessed (which path, which factorization, which deformation), you are no longer doing posets. You are flirting with genuinely higher structure: multiple morphisms between the same endpoints, and potential equivalences between those morphisms. That is exactly the territory where higher categories and HoTT naturally live. But there is a risk: it is easy to overclaim that polyhedral duality "is" HoTT or that HoTT is "needed." It is not. The right stance is restrained: the incidence category viewpoint is the base layer; higher-categorical and HoTT ideas become relevant only if you want to remember and reason about witness structure, coherence, and homotopy-like identification of paths.

10.1 From incidence categories to cell complexes and globular shapes

A convex polyhedron can be viewed as a 2-dimensional cell complex embedded in R^3: vertices (0-cells), edges (1-cells), and polygonal faces (2-cells) glued along boundaries. In combinatorial topology, one often abstracts away the embedding and keeps only the gluing data. The face poset is one way to record that, but it is not the only one.

If you think categorically, a cell complex suggests multiple layers of "composition":

- edges compose along vertices (paths in the 1-skeleton),
- faces glue along edge cycles (2-dimensional attachment),
- and in higher-dimensional complexes, higher cells attach along spheres.

The incidence category I(P) we used is intentionally minimal: it records inclusion of cells, not path concatenation or gluing maps. It is therefore a *shadow* of the full topological structure.

Higher categories enter when we try to represent not just objects and arrows, but also 2-dimensional and higher-dimensional composition laws. There are different "shapes" of higher categories depending on what you want composition to mean. A common one is **globular** shape: 0-morphisms, 1-morphisms between them, 2-morphisms between 1-morphisms, etc., arranged like nested source/target boundaries. Another is **simplicial** or **cubical** shape, where higher cells look like simplices or cubes and compositions correspond to gluing along faces.

Polyhedra and polytopes sit at a natural intersection of these perspectives:

- Their combinatorics is often expressed by face lattices and polytopal complexes (simplicial/cubical flavor).
- Their "directed" structure can be expressed by choosing orientations and thinking of boundary maps (globular flavor, if you encode cells as morphisms).

A useful conceptual bridge is this: once you have a polyhedron as a cell complex, you can consider its **fundamental groupoid** (or higher groupoid) which is inherently homotopical. That groupoid is not the same as the incidence category. But the incidence category can be seen as a scaffold from which richer constructions (like nerves, classifying spaces, or higher groupoids) can be built if desired.

So, “from incidence categories to cell complexes” means:

- Incidence categories encode “which cell is a face of which.”
- Cell-complex structures encode “how boundaries attach and how paths compose.”
- Higher categories provide formal languages for tracking compositions beyond the 1-dimensional incidence arrows.

10.2 Where “higher morphisms” could live: homotopies between incidence paths

In the thin/poset setting, between two objects there is at most one morphism. That kills witness structure and therefore kills the need for higher morphisms. But recall the alternative we discussed: the free category on the incidence graph. In that free category:

- morphisms are incidence paths,
- there can be many distinct morphisms v -> f (different v -> e -> f paths via different edges),
- and there can be many distinct morphisms between other endpoint pairs as well.

Now a new question becomes meaningful:

- when should two different incidence paths be considered “the same in a higher sense”?

This is precisely where 2-morphisms could live: as identifications (or transformations) between parallel 1-morphisms.

In a polyhedron, there is an intuitive source of such identifications. Consider a fixed face f of a polyhedron. Inside that face, there are multiple boundary edges meeting at each vertex. Two different incidence paths from a vertex to the face might correspond to choosing different boundary edges. But from the standpoint of “the vertex lies on the face,” those are the same fact witnessed in different ways. In a higher-categorical view, you might want:

- a 2-morphism between the two distinct 1-morphisms v -> f arising from different boundary edges.

More generally, if you build a category whose morphisms are combinatorial paths in a graph (like the 1-skeleton), then homotopies of paths in the cell complex can be used to generate 2-morphisms. In a 2-dimensional complex, a face provides a basic homotopy move: you can slide a path across a face, replacing one boundary route with another. This is a classical topological idea:

- 2-cells witness relations among 1-dimensional paths.

Translated into higher-category language:

- 1-morphisms are paths,
- 2-morphisms are homotopies between paths,
- faces generate such homotopies by filling loops.

In HoTT, this idea appears as the hierarchy:

- terms (points),
- identity types (paths between points),
- higher identity types (homotopies between paths),
- and so on.

So if we choose to interpret certain combinatorial objects (like paths in a skeleton) as "identity proofs," then 2-cells naturally become higher proofs of equality between those paths.

However, caution is needed: incidence paths v -> e -> f are not the same thing as paths in the 1-skeleton of the polyhedron. They are paths in the incidence graph of cells. That incidence graph is bipartite/tripartite by dimension and has a different interpretation. Still, the meta-idea remains: once you allow multiple morphisms between the same endpoints, you can ask for higher structure that identifies them under some equivalence relation, and that equivalence relation can be inspired by topological deformation or by combinatorial rewrite rules.

Thus, a plausible location for "higher morphisms" in this story is:

- start with the free incidence category,
- add 2-morphisms that identify incidence paths that represent the same containment fact,
- impose coherence conditions so these identifications behave consistently.

Whether that is valuable depends on what you want to measure or visualize. It is not needed to state duality, but it can be a meaningful enrichment if your goal is to treat the polyhedron not just as a poset but as a source of higher structure.

10.3 A restrained proposal: what HoTT adds without overclaiming

Here is the restrained proposal that fits the spirit of this paper:

1. **Base layer (what we already did):** treat a polyhedron P as a thin incidence category (face poset category). Duality is then an order-reversing equivalence, i.e. an equivalence

to an opposite category. This captures the combinatorial essence of duality and is sufficient for the cube–octahedron figure.

2. **Optional enrichment (where higher ideas enter):** if you want the diagram to carry more meaning than yes/no incidence—specifically, if you want to retain witness paths and reason about their equivalences—then you can replace the thin category by a richer structure (free category, or a 2-category, or a groupoid-like object). In that richer structure, faces (2-cells) can be used to generate 2-morphisms that identify different witness paths.

3. **HoTT's contribution:** HoTT provides a disciplined language for "witnesses and identifications of witnesses," and it forces you to be explicit about levels:

    - a claim of incidence can be treated as a proposition with possible proofs (witnesses),
    - equality of proofs (or equivalence of paths) becomes a higher-level object,
    - and coherence constraints become first-class rather than handwaved.

4. **What we should not claim:** we should not claim that duality "requires" HoTT, or that our figure is secretly a HoTT theorem. Duality already lives comfortably in posets and thin categories. HoTT becomes relevant only if we enrich the structure to track witness multiplicity and homotopy-like identifications.

5. **A concrete payoff (if one wants it):** a HoTT/higher-category enrichment could support new kinds of graphics:

    - not only showing which incidences exist,
    - but showing families of equivalent incidence proofs,
    - and showing how 2-cells mediate rewrite rules between incidence paths.

In other words, HoTT can add a second dimension of meaning: it can turn "there exists an incidence arrow" into "there exists a space (possibly contractible, possibly not) of incidence witnesses," and it can ask whether that witness space is trivial or has structure. That is a genuinely novel direction for "category pictures" of polyhedra if you pursue it carefully.

But the disciplined conclusion is this: the cube–octahedron duality figure is already a categorical object at the 1-category level. HoTT is not the foundation of that claim; it is a possible next layer if one decides to treat incidence not as a collapsed order relation but as a witness-rich system where "sameness" itself has structure.

11. Graphics as Proof Objects: Why This Visualization Matters

A reader might reasonably ask: why do any of this visually? If duality is a theorem about face posets and opposite categories, why not simply state it and move on? The answer is that, in practice, many duality claims drift into slogan form. People say "faces become vertices" and leave the rest to intuition. The visualization matters because it makes the claim inspectable. It is a page-sized object that can function like an audit trail: it shows exactly what has been asserted, how the asserted correspondences line up with incidence patterns, and where a purported duality could fail. In that sense, the graphic behaves like a proof object—not a proof by itself, but a structured witness whose correctness is checkable and whose mistakes are discoverable.

The deeper payoff is pedagogical and methodological. The diagram helps move the reader from "I can picture a cube and an octahedron" to "I can reason about invariants under duality in a language that scales." The aesthetic payoff is not accidental either: the diagram compresses a lot of structure into a single page in a way that remains truthful to the mathematics. This section explains why that is a meaningful achievement.

11.1 The diagram as an audit trail for duality claims

The strongest practical use of the diagram is not that it *convinces* you duality is true, but that it tells you precisely what must be true if the duality claim is correct.

A duality claim has three common failure modes:

- the object matching is wrong (faces matched to the wrong vertices, edges mismatched, etc.),
- the object matching is right but the incidence preservation is wrong (adjacencies do not correspond),
- the incidence preservation is right locally but fails globally (composites do not correspond coherently).

A purely verbal explanation often hides these failures. A diagram can expose them.

In the figure, each panel is already an "incidence audit" for one polyhedron: you can check counts, degrees, and boundary relations by tracing lines. Then the dotted correspondences provide a cross-panel audit handle: they fix the intended identification between specific named objects. Because the objects are labeled in rigid coordinate language (F_{+x}, V_{-z}, sign triples, etc.), the diagram prevents a common kind of vagueness: it forces the duality map to have a specific meaning, not just "up to symmetry."

Once those identifications are fixed, the audit becomes mechanical:

- a cube face has 4 boundary edges; its dual octa vertex must have 4 incident edges,
- a cube vertex has 3 incident edges; its dual octa face must have 3 boundary edges,
- a cube edge lies on 2 faces; its dual octa edge must have 2 endpoints (always) and lie on 2 faces, matching the “2-and-2” incidence.

These are not “proofs,” but they are necessary constraints. If the diagram violated any of them, it would be wrong in a way that is immediately visible.

More subtly, the diagram suggests checks that correspond to functorial integrity:

- whenever v -> e -> f exists on the cube side (i.e., a vertex on a boundary edge of a face), the dual should exhibit d(f) -> d(e) -> d(v) on the octa side.

Even without drawing every cross-panel correspondence, the figure’s labeling scheme makes it possible to reconstruct those checks. That is what makes it an audit trail: it encodes enough information that a reader can verify the intended structure without guessing what the author meant.

So the diagram is “proof-like” in the sense that it is falsifiable. It is an object you can interrogate. If the author has smuggled in a wrong correspondence, the diagram gives you the machinery to detect it.

11.2 Pedagogy: from geometric intuition to categorical invariants

There is a predictable learning curve with duality. Most people first meet it as a geometric trick: “put a point in each face and connect points across edges.” That is good intuition, but it is not a stable language for generalization. It is hard to scale that intuition:

- to polyhedra with uneven faces,
- to higher dimensions,
- to computational settings where you do not want to rely on spatial reasoning.

The category/poset viewpoint solves this by shifting the object of study. Instead of studying the polyhedron as a shape in space, you study its face-incidence structure as a categorical invariant. This viewpoint has several pedagogical advantages that the diagram reinforces.

It clarifies what duality preserves and what it reverses:

- It preserves “the existence of incidence chains” up to reversal.
- It reverses the partial order of face inclusion.

Once you see the diagram as two incidence categories with an implicit op relationship, the phrase "duality is contravariant" stops being abstract. It becomes a literal reading rule: arrows reverse after swapping strata.

It also helps separate the roles of different kinds of information:

- metric information (lengths, angles) is absent,
- combinatorial information (which touches which) is present and central,
- symmetry is optional and not required for the categorical statement.

This is pedagogically important because it teaches the reader what is essential. The cube–octahedron pair is symmetric, but duality does not depend on symmetry. The diagram uses symmetry to make the presentation clean, but the category-theoretic lesson is that the same structure persists without it.

Finally, the diagram hints at a broader conceptual move: invariants under equivalence. Once you view the face poset as a category, you can ask what is invariant under categorical equivalence (or under isomorphism in the thin setting). In plain terms:

- which properties are really properties of the incidence category itself?

This is the bridge from geometric intuition to categorical invariants. The diagram is a stepping stone that lets a reader carry their spatial intuition into a more abstract but more general framework.

11.3 Aesthetic compression: fitting structure onto a page without lying

Mathematical graphics have a moral hazard: they can oversimplify and thereby mislead. The aim here is the opposite: to compress structure without distorting it. The figure's aesthetic success comes from three design constraints that are also mathematical constraints.

First, it uses stratification by dimension. This prevents a common confusion in polyhedral diagrams where vertices, edges, and faces are mixed together in a way that hides what is an object and what is a relation. Here, each stratum is explicitly "objects of one rank." The lines are therefore unambiguously "incidence morphisms," not geometric edges drawn to look like 3D solids.

Second, it uses minimal generating relations. Only V -> E and E -> F are drawn. This is an honest choice because the rest (like V -> F) is categorically forced by composition in the thin reading. The diagram is not hiding information; it is showing the generators and relying on categorical closure as a legitimate compression method.

Third, it uses rigid labels rather than freehand geometry. The diagram does not ask the reader to "recognize" a cube or an octahedron by shape. It asks the reader to recognize a structure by names and incidence. This eliminates the temptation to read the picture as perspective art. It is a labeled graph representing a category, not a drawing representing a solid.

These constraints produce an aesthetic that is not ornamental but structural: the beauty comes from seeing a large amount of correct structure packed into a small space. The diagram "fits on a page" because it is carefully choosing what level of abstraction to draw. It does not attempt to draw all morphisms (which would clutter it), and it does not attempt to draw an embedding in 3D (which would distract from the categorical claim). It draws exactly what the thesis needs.

So the visualization matters because it is:

- falsifiable (audit-friendly),
- pedagogically efficient (it teaches the categorical view),
- and aesthetically compressed without losing truth (it shows generators and lets composition do the rest).

That is a respectable role for a mathematical graphic: it is not merely illustration; it is a structured witness to a claim about equivalence of incidence data under duality.

12. Conclusion and Forward Work

The cube–octahedron figure began as a visual curiosity—two Platonic solids with swapped counts—and ended as a categorical thesis: duality is best understood as a structured reversal of incidence, formalized as an equivalence to an opposite category. That framing matters because it stabilizes what people often treat as a geometric trick into something that is definitionally checkable. When the duality is written as a relationship between incidence categories (or face-poset categories), it is no longer "an intuition." It becomes a claim with an audit trail: a specific object mapping, a specific direction reversal, and specific commuting conditions. The figure is valuable precisely because it makes those requirements visible while remaining small enough to fit on a page.

12.1 Duality restated as structured reversal, not a trick

The headline restatement is simple:

- polyhedral duality is an order-reversing correspondence on faces,
- and order reversal is the categorical operation "take the opposite."

In the thin/poset formulation, the entire story can be compressed into one sentence:

- $F(P^*) \simeq F(P)^{op}$.

Everything else is unpacking what that sentence means in a way that a reader can see. In the three-stratum diagram, "op" becomes a rule:

- swap V with F, keep E,
- reverse incidence direction.

This reframes duality away from construction-specific language ("put points in faces and connect them") and toward structure ("incidence data reverses"). The duality becomes a map that is forced to respect the logic of containment. That is why the figure can be treated as a proof-adjacent object: it does not just show the correspondence; it encodes the directionality that makes the correspondence meaningful.

In that sense, duality is not a trick. It is a functorial phenomenon: a contravariant equivalence of incidence structures.

12.2 Next figures: tetrahedron self-duality; dodecahedron–icosahedron

The cube–octahedron figure is one member of a family of "category pictures" for Platonic dual pairs. Two immediate next figures suggest themselves, each with a different pedagogical role.

The tetrahedron (self-duality). The tetrahedron is the minimal nontrivial case and the cleanest demonstration of what "self-dual" means categorically. In 3D, self-duality amounts to an isomorphism:

- $F(T) \simeq F(T)^{op}$

(or, more carefully, $T^* \cong T$ and hence $F(T^*) \simeq F(T)^{op}$, which then identifies with $F(T)$ via the self-duality). A category picture for the tetrahedron would show that:

- 4 vertices $\leftrightarrow$ 4 faces,
- 6 edges $\leftrightarrow$ 6 edges,

with the incidence reversal happening "inside one object." This is the place where the opposite-category principle becomes most visually striking: the same incidence diagram, read upside down (after relabeling), gives you the same structure. It also illustrates a subtle point: "self-dual" is a property of the incidence structure, not a property of a particular embedding.

The dodecahedron–icosahedron pair. This is the "large but still human" demonstration. Here the numbers are bigger:

- dodecahedron: 20 vertices, 30 edges, 12 faces,
- icosahedron: 12 vertices, 30 edges, 20 faces.

The duality again swaps V and F and fixes E. But unlike cube–octahedron, the faces are pentagons vs triangles, and the local incidence degrees differ correspondingly. A category picture for this pair would emphasize how duality exchanges "pentagonal boundary" with "vertex valence 5," and "triangular boundary" with "vertex valence 3." It would also stress that symmetry is helpful but not essential: even if you scrambled labels, the incidence category still determines the dual. The challenge (and opportunity) in this figure is aesthetic compression: how to draw 20–30–12 objects without creating noise. That challenge is exactly what motivates automation and invariants extraction.

12.3 Open directions: automated category pictures and invariants extraction

The most forward-looking direction is to treat "category pictures" as something you can generate automatically from incidence data, not as a one-off hand-crafted diagram. This points to a practical workflow:

1. Input: a polytope described combinatorially (lists of faces, or adjacency relations).
2. Build: the face poset (thin category) or an enriched incidence category (free/path-sensitive).
3. Compute: the dual by order reversal (op), optionally identify it with a known polytope.
4. Render: a stratified diagram (Hasse-style layers by rank) with stable labeling and controlled edge routing.

Once this is automated, you can go beyond "pretty pictures" and extract invariants directly from the incidence category. Examples of invariants that fall naturally out of the categorical/poset view include:

- rank counts: the f-vector (f_0, f_1, ..., f_{n-1}),
- degree distributions: vertex valences and face polygon sizes (dual to each other),
- Euler-type constraints in low dimensions,
- lattice/poset properties: gradedness, atomicity, modularity failures, etc.,
- symmetry signatures: automorphisms of the face poset (categorical automorphisms).

Even more interesting is the prospect of extracting *structural fingerprints* that are stable under isomorphism of incidence categories: canonical forms of posets, invariants derived from nerves, or summary statistics of the Hasse diagram. For large polytopes, these become the “compressed truth” of the object: a way to compare polytopes without drawing them in 3D.

A second open direction is controlled enrichment: deciding how much witness structure to keep. The thin category is canonical and clean, but sometimes you may want more:

- keep multiple incidence paths and study their equivalences,
- attach 2-cells that witness path rewrites,
- connect the incidence scaffold to fundamental groupoid/homotopy data of the cell complex.

This is where higher categories and HoTT can enter in a disciplined way, not as a replacement for the poset story but as an optional upgrade.

Finally, there is a “researcher’s payoff” direction: use these automated pictures as experimental probes. For example:

- generate category pictures for families of polytopes,
- compute invariants,
- cluster by incidence fingerprints,
- and search for surprising near-dualities, self-dualities, or symmetry-breaking transitions.

That kind of pipeline turns the aesthetic practice of drawing diagrams into a small empirical program: graphics as a means of discovery as well as explanation.

The conclusion, then, is not merely that the cube and octahedron are dual. It is that the right way to say “dual” in a scalable, auditable language is:

- duality is contravariant structure preservation,
- and the opposite-category operation is the native formal expression of incidence reversal.

The figure works because it embodies that claim in a form a reader can inspect. The forward work is to turn that embodiment into a repeatable instrument: automated category pictures, systematic duality checks, and invariants extraction that make “structure on a page” a reliable part of mathematical practice rather than a decorative afterthought.

References

[1] E. W. Weisstein, "Dual Polyhedron," *MathWorld—A Wolfram Web Resource*. MathWorld

[2] E. W. Weisstein, "Platonic Solid," *MathWorld—A Wolfram Web Resource*. MathWorld

[3] E. W. Weisstein, "Cube," *MathWorld—A Wolfram Web Resource*. MathWorld

[4] E. W. Weisstein, "Cube-Octahedron Compound," *MathWorld—A Wolfram Web Resource*. MathWorld

[5] P. McMullen and E. Schulte, *Abstract Regular Polytopes*, Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2002. doi:10.1017/CBO9780511546686. Cambridge University Press & Assessment+1

[6] T. Leinster, *Higher Operads, Higher Categories*, London Mathematical Society Lecture Note Series 298. Cambridge University Press, 2004. ISBN 9780521532150. Also available as arXiv:math/0305049. Cambridge Assets+2arXiv+2

[7] J. Lurie, *Higher Topos Theory*, Annals of Mathematics Studies 170. Princeton University Press, 2009. (Also circulated as preprint PDF.) math.ias.edu+2arXiv+2

[8] The Univalent Foundations Program (Institute for Advanced Study), *Homotopy Type Theory: Univalent Foundations of Mathematics*, 2013. Homotopy Type Theory+2Homotopy Type Theory+2

[9] A. A. Hagberg, D. A. Schult, and P. J. Swart, "Exploring Network Structure, Dynamics, and Function Using NetworkX," in *Proceedings of the 7th Python in Science Conference (SciPy 2008)*, 2008. proceedings.scipy.org+2aric.hagberg.org+2

[10] J. D. Hunter, "Matplotlib: A 2D Graphics Environment," *Computing in Science & Engineering*, 9(3):90–95, 2007. doi:10.1109/MCSE.2007.55. ACM Digital Library+1

[11] C. R. Harris et al., "Array programming with NumPy," *Nature* 585, 357–362, 2020. doi:10.1038/s41586-020-2649-2. Nature+1

[12] P. Virtanen et al., "SciPy 1.0: fundamental algorithms for scientific computing in Python," *Nature Methods* 17, 261–272, 2020. doi:10.1038/s41592-019-0686-2. Nature+1

[13] G. M. Ziegler, *Lectures on Polytopes*, Graduate Texts in Mathematics 152. Springer, 1995. Google Books

[14] B. Grünbaum, *Convex Polytopes*. Interscience (original ed.), 1967; widely reprinted and cited. Internet Archive

[15] A. Brøndsted, *An Introduction to Convex Polytopes*. Springer, 1983. SciSpace

[16] H. S. M. Coxeter, *Regular Polytopes*. Dover Publications, 1973 (Dover reprint/edition commonly cited). eBay

[17] M. J. Wenninger, *Dual Models*. Cambridge University Press, 1983. (Historical/polyhedral modeling perspective on duality; cited by MathWorld for early duality references.) MathWorld

[18] (Optional computational cross-check) Wolfram Demonstrations Project, “Duals of Platonic Solids.” Wolfram Demonstrations

[19] S. Mac Lane, *Categories for the Working Mathematician*, 2nd ed., Graduate Texts in Mathematics 5. Springer, 1998.

[20] T. Leinster, *Basic Category Theory*. Cambridge University Press, 2014.

[21] S. Awodey, *Category Theory*, 2nd ed. Oxford University Press, 2010.

[22] E. Riehl, *Category Theory in Context*. Dover Publications, 2016.

Python Source Code

```
"""
Duality as an incidence "category picture" (Cube <-> Octahedron)

Generates a figure with:
- I(C): cube incidence (V -> E solid, E -> F dashed)
- I(O): octa incidence (V -> E solid, E -> F dashed)
- Center correspondences: cube faces F(C) <-> octa vertices V(O) (black dotted)

Outputs:
- duality_incidence_mathlabels_compact_raised.png
- duality_incidence_mathlabels_compact_raised.svg
"""

from __future__ import annotations

from dataclasses import dataclass
from typing import Dict, List, Tuple, Iterable, FrozenSet
import numpy as np
import matplotlib.pyplot as plt


# -----------------------------
# Helpers / canonical ordering
# -----------------------------
```

```
def hamming_dist(a: Tuple[int, int, int], b: Tuple[int, int, int]) -> int:
    return sum(int(ai != bi) for ai, bi in zip(a, b))


def sign_str(s: int) -> str:
    return "+" if s > 0 else "-"


def cube_vertex_order() -> List[Tuple[int, int, int]]:
    """
    Matches the visual ordering used in the figure (top to bottom):
    (-1,-1,-1), (-1,-1,+1), (-1,+1,-1), (-1,+1,+1),
    (+1,-1,-1), (+1,-1,+1), (+1,+1,-1), (+1,+1,+1)
    """
    V = []
    for x in (-1, +1):
        for y in (-1, +1):
            for z in (-1, +1):
                V.append((x, y, z))
    # That loop yields exactly the ordering above.
    return V


def cube_edges(vertices: List[Tuple[int, int, int]]) -> List[Tuple[int, int]]:
    """
    Return cube edges as index pairs (i,j), sorted lex by (i,j),
    where i < j and vertices differ by Hamming distance 1.
    """
    idx = {v: k for k, v in enumerate(vertices)}
```

```
    E = []
    for i in range(len(vertices)):
        for j in range(i + 1, len(vertices)):
            if hamming_dist(vertices[i], vertices[j]) == 1:
                E.append((i, j))
    E.sort()
    return E


def cube_faces() -> List[Tuple[str, int]]:
    """
    Six faces labeled by coordinate and sign: (axis, sign)
    Order matches figure (top to bottom): +x, -x, +y, -y, +z, -z
    """
    return [("x", +1), ("x", -1), ("y", +1), ("y", -1), ("z", +1), ("z", -1)]


def face_label(axis: str, s: int) -> str:
    return f"F_{{{sign_str(s)}{axis}}}"


def octa_vertex_label(axis: str, s: int) -> str:
    return f"V_{{{sign_str(s)}{axis}}}"


def octa_face_label(sx: int, sy: int, sz: int) -> str:
    # Matches the figure's far-right labels: F_{+x+y+z}, etc.
    return f"F_{{{sign_str(sx)}x{sign_str(sy)}y{sign_str(sz)}z}}"
```

```
# -----------------------------
# Build combinatorics
# -----------------------------

# Cube
V_C = cube_vertex_order()
E_C = cube_edges(V_C)
F_C = cube_faces()  # (axis, sign)

# Face membership for cube vertices
def cube_vertex_in_face(v: Tuple[int, int, int], face: Tuple[str, int]) -> bool:
    axis, s = face
    x, y, z = v
    if axis == "x":
        return x == s
    if axis == "y":
        return y == s
    if axis == "z":
        return z == s
    raise ValueError("bad axis")

# Edge -> incident vertices
edge_to_vertices_C: Dict[int, Tuple[Tuple[int, int, int], Tuple[int, int, int]]] = {}
for e_id, (i, j) in enumerate(E_C):
    edge_to_vertices_C[e_id] = (V_C[i], V_C[j])
```

```
# Cube edge -> incident faces (exactly two)
edge_to_faces_C: Dict[int, List[Tuple[str, int]]] = {}
for e_id, (v1, v2) in edge_to_vertices_C.items():
    incident = []
    for f in F_C:
        if cube_vertex_in_face(v1, f) and cube_vertex_in_face(v2, f):
            incident.append(f)
    # Each cube edge lies on exactly two faces
    edge_to_faces_C[e_id] = incident

# Octahedron vertices are dual to cube faces
V_O = F_C[:]  # represent octa vertices by the same (axis, sign) pairs

# Duality component shown in the figure:
# d_F : F(C) -> V(O)  (identity on our representation)
def d_face_to_octa_vertex(face: Tuple[str, int]) -> Tuple[str, int]:
    return face

# Octa edges are dual to cube edges via "edge = intersection of two faces"
# For each cube edge e with faces f1,f2, define octa edge connecting d(f1), d(f2)
OctaEdge = FrozenSet[Tuple[str, int]]
E_O: List[OctaEdge] = []
edgeC_to_edgeO: Dict[int, OctaEdge] = {}
for e_id in range(len(E_C)):
    f1, f2 = edge_to_faces_C[e_id]
    oe = frozenset([d_face_to_octa_vertex(f1), d_face_to_octa_vertex(f2)])
```

```
    edgeC_to_edgeO[e_id] = oe
    E_O.append(oe)

# Make E_O unique and preserve cube-edge order (it already is a bijection).
# We'll use cube edge indices as canonical edge labels E0..E11 in both panels.

# Octa faces are dual to cube vertices (sign triples)
# Represent octa faces as sign triples (sx,sy,sz) and their 3 vertices:
# { (x,sx), (y,sy), (z,sz) }
def octa_face_vertices(sx: int, sy: int, sz: int) -> FrozenSet[Tuple[str, int]]:
    return frozenset([("x", sx), ("y", sy), ("z", sz)])

# Order for octa faces (top to bottom) to match the figure:
# +++ , ++- , +-+ , +-- , -++ , -+- , --+ , ---
octa_face_sign_order = [
    (+1, +1, +1),
    (+1, +1, -1),
    (+1, -1, +1),
    (+1, -1, -1),
    (-1, +1, +1),
    (-1, +1, -1),
    (-1, -1, +1),
    (-1, -1, -1),
]
F_O: List[Tuple[int, int, int]] = octa_face_sign_order[:]
faceO_to_vertices: Dict[int, FrozenSet[Tuple[str, int]]] = {
```

```
    f_id: octa_face_vertices(*signs) for f_id, signs in enumerate(F_O)
}

# Octa edge membership: an octa edge is a 2-subset of vertices that share a face.
# We already have E_O from duality; build incidence E->F for octa by testing inclusion.
edgeO_to_facesO: Dict[int, List[int]] = {e_id: [] for e_id in range(len(E_O))}
for e_id, oe in enumerate(E_O):
    for f_id, vset in faceO_to_vertices.items():
        if oe.issubset(vset):
            edgeO_to_facesO[e_id].append(f_id)
    # Each octa edge lies on exactly 2 triangular faces
    # (We won't assert hard, but it should be length 2 in a correct build)



# -----------------------------
# Layout (page-fit / "compact")
# -----------------------------

# Normalized coordinates in [0,1] (axis turned off)
x_cube_V = 0.08
x_cube_E = 0.24
x_cube_F = 0.39

x_octa_V = 0.54
x_octa_E = 0.70
x_octa_F = 0.90
```

```
top = 0.86
bottom = 0.10

# Vertical placements
y_cube_V = np.linspace(top, bottom, len(V_C))        # 8 cube vertices
y_edges = np.linspace(top, bottom, len(E_C))         # 12 edges (shared y for both panels)
y_mid = np.linspace(top, bottom, len(F_C))           # 6 cube faces / octa vertices
y_octa_F = np.linspace(top, bottom, len(F_O))        # 8 octa faces

# Label offsets
dx_left_label = 0.015
dx_right_label = 0.012
label_raise = 0.012  # <-- "raise text slightly" above the center dotted lines



# -----------------------------
# Drawing primitives
# -----------------------------

def draw_line(ax, p, q, **kw):
   ax.plot([p[0], q[0]], [p[1], q[1]], **kw)

def draw_node(ax, x, y, marker, size, facecolor, edgecolor="none", zorder=3):
   ax.scatter([x], [y], s=size, marker=marker, c=[facecolor], edgecolors=edgecolor, zorder=zorder)
```

```
# -----------------------------
# Render
# -----------------------------

def render(out_png: str, out_svg: str) -> None:
    fig = plt.figure(figsize=(8.6, 8.6), dpi=220)
    ax = fig.add_axes([0, 0, 1, 1])
    ax.set_xlim(0, 1)
    ax.set_ylim(0, 1)
    ax.axis("off")

    # Titles
    ax.text(0.5, 0.965, "Duality:  d : F(C) <-> V(O)", ha="center", va="center", fontsize=14)
    ax.text(0.22, 0.90, "I(C)", ha="center", va="center", fontsize=12)
    ax.text(0.73, 0.90, "I(O)", ha="center", va="center", fontsize=12)

    # --- Cube panel: vertices
    for i, v in enumerate(V_C):
        y = float(y_cube_V[i])
        draw_node(ax, x_cube_V, y, marker="o", size=90, facecolor="tab:blue")
        ax.text(x_cube_V - dx_left_label, y, f"({v[0]},{v[1]},{v[2]})", ha="right", va="center",
fontsize=10)

    # --- Edge nodes (shared y for cube and octa), label as E0..E11
    for e_id in range(len(E_C)):
```

```
        y = float(y_edges[e_id])
        # cube edge objects
        draw_node(ax, x_cube_E, y, marker="s", size=70, facecolor="tab:orange")
        ax.text(x_cube_E, y, f"E{e_id}", ha="center", va="center", fontsize=8)

        # octa edge objects
        draw_node(ax, x_octa_E, y, marker="s", size=70, facecolor="tab:orange")
        ax.text(x_octa_E, y, f"E{e_id}", ha="center", va="center", fontsize=8)

    # --- Cube face objects in the center (green triangles)
    for f_id, f in enumerate(F_C):
        y = float(y_mid[f_id])
        draw_node(ax, x_cube_F, y, marker="^", size=85, facecolor="tab:green")

        # label raised above the center dotted line
        ax.text(x_cube_F + 0.010, y + label_raise, face_label(*f), ha="left", va="center",
fontsize=10)

    # --- Octa vertices (blue dots) aligned with cube faces
    for v_id, ov in enumerate(V_O):
        y = float(y_mid[v_id])
        draw_node(ax, x_octa_V, y, marker="o", size=90, facecolor="tab:blue")

        ax.text(x_octa_V - 0.010, y + label_raise, octa_vertex_label(*ov), ha="right", va="center",
fontsize=10)

    # --- Octa faces on far right (green triangles) with sign-triple labels
```

```
    for f_id, (sx, sy, sz) in enumerate(F_O):
        y = float(y_octa_F[f_id])
        draw_node(ax, x_octa_F, y, marker="^", size=85, facecolor="tab:green")
        ax.text(x_octa_F + dx_right_label, y, octa_face_label(sx, sy, sz), ha="left", va="center",
fontsize=10)

    # --- Center correspondences: cube faces <-> octa vertices (black dotted)
    for k, f in enumerate(F_C):
        y = float(y_mid[k])
        draw_line(ax, (x_cube_F, y), (x_octa_V, y), color="black", linestyle=":", linewidth=1.2,
zorder=1)

    # --- Incidence edges inside cube panel
    # V(C) -> E(C): each cube edge has 2 endpoints
    for e_id, (i, j) in enumerate(E_C):
        y_e = float(y_edges[e_id])
        y_v1 = float(y_cube_V[i])
        y_v2 = float(y_cube_V[j])
        draw_line(ax, (x_cube_V, y_v1), (x_cube_E, y_e), color="tab:blue", linewidth=1.0, zorder=1)
        draw_line(ax, (x_cube_V, y_v2), (x_cube_E, y_e), color="tab:blue", linewidth=1.0, zorder=1)

    # E(C) -> F(C): each cube edge lies on 2 faces
    face_index_C = {f: k for k, f in enumerate(F_C)}
    for e_id in range(len(E_C)):
        y_e = float(y_edges[e_id])
        for f in edge_to_faces_C[e_id]:
            y_f = float(y_mid[face_index_C[f]])
```

```
    draw_line(
        ax,
        (x_cube_E, y_e),
        (x_cube_F, y_f),
        color="tab:blue",
        linestyle="--",
        linewidth=1.0,
        zorder=1,
    )

# --- Incidence edges inside octa panel
# V(O) -> E(O): each octa edge has 2 endpoints (vertices are in V_O)
octa_vertex_index = {v: k for k, v in enumerate(V_O)}
for e_id, oe in enumerate(E_O):
    y_e = float(y_edges[e_id])
    a, b = tuple(oe)
    y_a = float(y_mid[octa_vertex_index[a]])
    y_b = float(y_mid[octa_vertex_index[b]])
    draw_line(ax, (x_octa_V, y_a), (x_octa_E, y_e), color="tab:blue", linewidth=1.0, zorder=1)
    draw_line(ax, (x_octa_V, y_b), (x_octa_E, y_e), color="tab:blue", linewidth=1.0, zorder=1)

# E(O) -> F(O): each octa edge lies on 2 triangular faces
for e_id, face_ids in edgeO_to_facesO.items():
    y_e = float(y_edges[e_id])
    for f_id in face_ids:
        y_f = float(y_octa_F[f_id])
```

```
        draw_line(
            ax,
            (x_octa_E, y_e),
            (x_octa_F, y_f),
            color="tab:blue",
            linestyle="--",
            linewidth=1.0,
            zorder=1,
        )

    # Save
    fig.savefig(out_png, bbox_inches="tight", pad_inches=0.02)
    fig.savefig(out_svg, bbox_inches="tight", pad_inches=0.02)
    plt.close(fig)


if __name__ == "__main__":
    render(
        out_png="duality_incidence_mathlabels_compact_raised.png",
        out_svg="duality_incidence_mathlabels_compact_raised.svg",
    )
    print("Wrote duality_incidence_mathlabels_compact_raised.png and .svg")
```