Category theory and diagrammatic reasoning

6th February 2019

Last updated: 5th February 2019

2 Bicategories and natural transformations

In the first lecture, we saw how a composition operation defining a category over a reflexive graph determines a notion of *commutative diagram.* A converse is also true: from the data of the underlying reflexive graph X and all commutative diagrams $\partial O^2_{n,m} \to X$, we can reconstruct the composition operation. Indeed, for each path $(f_1, \ldots, f_n)$, there is going to be a unique f such that

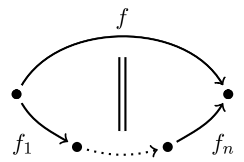

,

and that is the composite $f_1; \ldots; f_n$.

Thus we can think of presenting X as a reflexive graph together with an equivalence relation on the parallel finite paths in X. To be the equivalence relation arising from a composition, this needs to satisfy some properties:

- for each path $(f_1, \ldots, f_n)$, there is a unique edge f such that $(f_1, \ldots, f_n) \sim (f)$;
- for all $f : x \to y$, we have $(f, \mathrm{id}_y) \sim (f) \sim (\mathrm{id}_x, f)$;
- the equivalence relation "composes" as in [Lecture 1, Proposition 30]: that is, given any of the diagrams

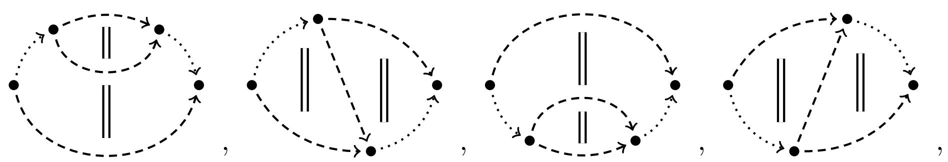

  the overall lower boundary is equivalent to the overall upper boundary.

In Lecture 1, we got to the concept of category starting from a graph presenting an equivalence relation, dropping the symmetry of the relation and turning "transitivity" into "composition": we can do the same here! First, we can present this equivalence relation with a graph whose vertices are the finite paths in another graph.

Definition 1. A (regular) *2-polygraph* X is a pair of a graph

$$X_1 \underset{t}{\overset{s}{\rightrightarrows}} X_0$$

and a graph

$$X_2 \underset{t}{\overset{s}{\rightrightarrows}} X_1^+$$

such that $s; s = t; s$ and $s; t = t; t$.

A *reflexive 2-polygraph* is a 2-polygraph with the additional data of functions $\mathrm{id}_{(-)} : X_0 \to X_1$ and $\mathrm{id}_{(-)} : X_1^+ \to X_2$ which make the two graphs reflexive.

We will call the elements of X_0 the *0-cells*, the elements of X_1 the *1-cells*, and the elements of X_2 the *2-cells* of the 2-polygraph. We represent a 2-cell α whose source is the path $(f_1, \ldots, f_n)$ and target is the path $(g_1, \ldots, g_m)$ in the following way:

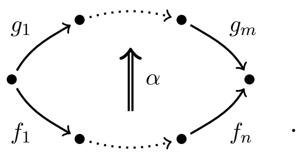

We write $\alpha : (f_1, \ldots, f_n) \Rightarrow (g_1, \ldots, g_m)$ for such a 2-cell. The conditions $s; s = t; s$ and $s; t = t; t$ stipulate that there can be a 2-cell only between two *parallel* paths.

Remark 2. Just like a graph can be seen as an abstract rewriting system, a 2-polygraph with a single 0-cell can be seen as a *string rewriting system*:

- the set of edges X_1, all necessarily going from and to the unique vertex, is an alphabet Σ;
- finite paths in the graph $s, t : \Sigma \to \{*\}$ are words of symbols in Σ;
- a 2-cell between two finite paths is a rewrite rule, saying that a word $a_1 \ldots a_n$ can be rewritten to a word $b_1 \ldots b_m$.

Definition 3. Let X be a reflexive 2-polygraph. A *composition* on X is an operation assigning to each diagram of 2-cells with shape

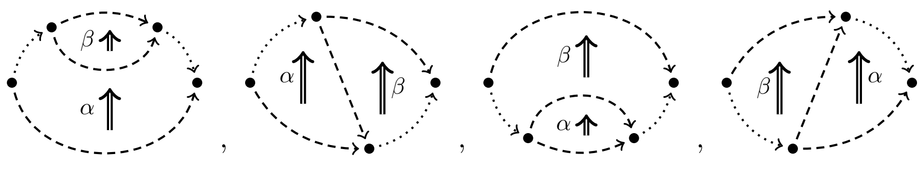

a unique composite, whose source and target are the overall lower and upper boundary of the diagram, in such a way that

1. composition is associative, in the sense that given a diagram of three 2-cells, if one can compose them pairwise in two different orders, the result is independent of the order;
2. for all paths p, the identity 2-cell id_p is a unit for composition, in the sense that

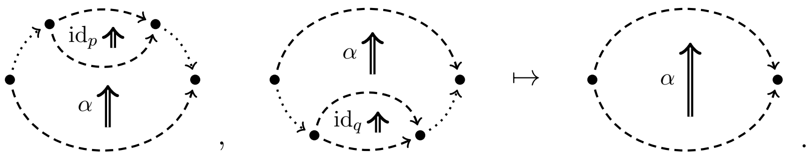

We will write, for example,

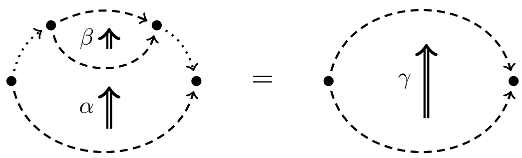

to mean that the right-hand side is the composite of the left-hand side. Diagrams like the ones we are drawing are often called *pasting diagrams*.

Remark 4. While it is easy to understand composition in pictures, the formalisation is more tricky. This is only a sample of possible configurations of three 2-cells:

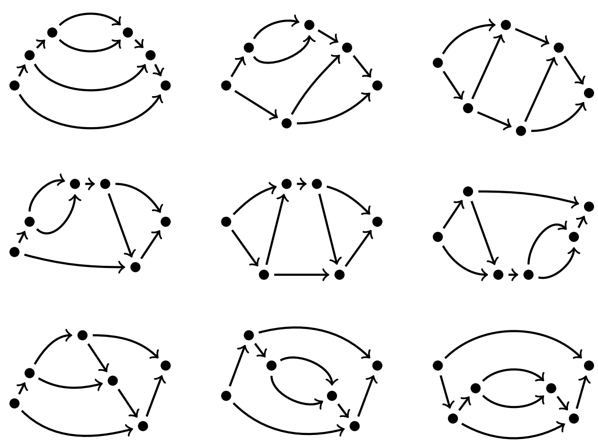

and we need an associativity axiom scheme for each one of these. Overall there are 33 different schemes! In any case, the explanation "in words" is good enough for practical purposes: it should suffice to know that it can be formalised.

Exercise 5. Let x, y be two 0-cells in a 2-polygraph with composition X. Show that there is a category $\mathrm{Hom}_X(x, y)$ whose objects are the 1-cells $f : x \to y$, and morphisms from $f : x \to y$ to $g : x \to y$ are the 2-cells $\alpha : (f) \Rightarrow (g)$.

For this definition to generalise what we were talking about — a category presented by “what its commutative diagrams are” — we need some special 2-cells to exhibit a 1-cell as a composite of other 1-cells, in such a way that the identity 1-cells are units. Recall that an equivalence relation, as a category, is the same thing as a preorder whose morphisms are all isomorphisms, so it makes sense to use 2-cells that are invertible.

Definition 6. Let X be a 2-polygraph with composition. A 2-cell $\alpha : q \Rightarrow p$ of X is an *isomorphism* (or *invertible*) if there is a 2-cell $\alpha^{-1} : q \Rightarrow p$ such that

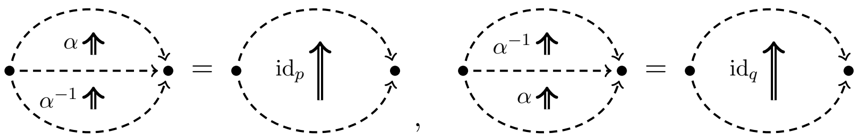

Definition 7. A *bicategory* is a 2-polygraph with composition together with:

1. for each pair $f : x \to y$ and $g : y \to z$ of 1-cells, a 1-cell $f; g : x \to z$ and an invertible 2-cell (the *compositor*)

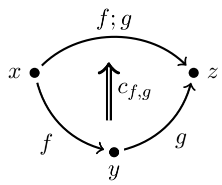

;

2. for each 1-cell $f : x \to y$, two invertible 2-cells (the left and right *unitors*)

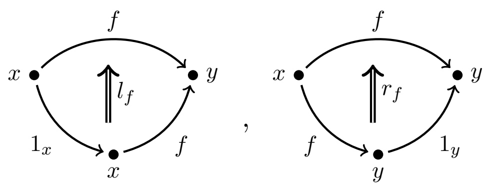

,

satisfying the following equations, for all compatible α, α', β, as pictured:

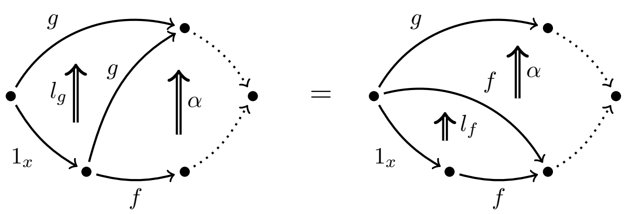

(1)

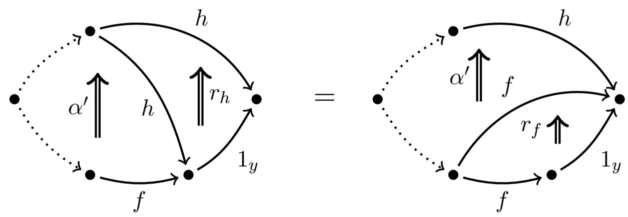

(2)

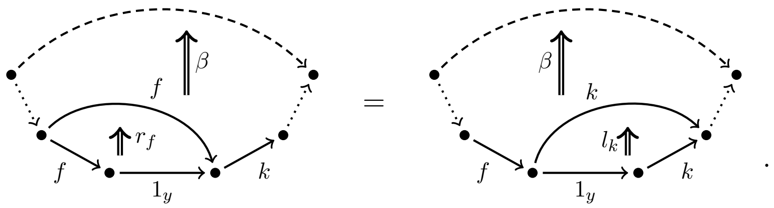

(3)

Remark 8. This is not the usual definition of bicategory that you will find in the literature; however, it is equivalent to it, and more directly useful for diagrammatic reasoning. Later in the course, we will also look at the "standard" definition in a special case, as it may be easier to check in practice.

The unitors exhibit the fact that an identity can be introduced to the left or to the right of any 1-cell. The equations (1—3) say that only the "position" of the identity matters, and not the way it was introduced: if there are two ways of applying unitors to a 2-cell, with the identity in the same place, then they have the same composite.

Exercise 9. Show that "duals" of the equations (1—3) hold with the source and target of each 2-cell reversed, using the inverses l_f^{-1} and r_f^{-1} instead of l_f and r_f, for each f.

Remark 10. The compositors induce a bijection between 2-cells with many inputs and many outputs, and 2-cells with one input and one output. Given a 2-cell $\alpha : (f_1, \ldots, f_n) \Rightarrow (g_1, \ldots, g_m)$, we can start post-composing it with compositors, and pre-composing it with their inverses. Depending on the order that we do it, we will get a single-input, single-output 2-cell between "iterated binary composites" with different bracketings.

For example, 2-cells $(f, g, h) \Rightarrow (k)$ corresponds bijectively to 2-cells

$$((f; g); h) \Rightarrow (k) \qquad \text{and} \qquad (f; (g; h)) \Rightarrow (k).$$

There is no reason that $(f; g); h$ and $f; (g; h)$ be *equal* as 1-cells: however, there will be an isomorphism between them.

While bicategories are fairly easy to find once you know about them, the following special case is more immediately recognisable in its ubiquity.

Definition 11. A *monoidal category* is a bicategory with a single 0-cell.

Whenever in a category there is a way of *composing objects* in an associative way, chances are the category is in fact a monoidal category:

- the 1-cells are the *objects* of the category, so finite paths of 1-cells are finite lists of objects;
- the 2-cells $(x_1, \ldots, x_n) \Rightarrow (y_1, \ldots, y_m)$ are the *morphisms* from the composite of the objects $(x_1, \ldots, x_n)$ to the composite of the objects $(y_1, \ldots, y_m)$.

Example 12. There is a monoidal category $\mathbf{Set}_\times$ whose 1-cells are sets, and a 2-cell $f : (S_1, \ldots, S_n) \Rightarrow (T_1, \ldots, T_m)$ is a function of n variables $(x_1, \ldots, x_n)$, with $x_i \in S_i$, returning an m-uple of values $(y_1, \ldots, y_m)$ with $y_i \in T_i$:

- identity 2-cells are identity functions, and the identity 1-cell is the set $\{*\}$ of a single element;

- composition of 2-cells amounts to substitution of a tuple of variables with the output of a function;

- for each pair (S, T) of sets, we let their composite be the cartesian product $S \times T$ of sets: then $c_{S,T} : (S, T) \to (S \times T)$ maps the pair (x, y) of an element of S and an element of T to the pair (x, y) as an element of the cartesian product $S \times T$;

- the unitor $l_S : (\{*\}, S) \to (S)$ maps the pair $(*, x)$ to x, and similarly $r_S : (S, \{*\}) \to S$ maps $(x, *)$ to x, for each $x \in S$.

Example 13. Suppose P is a meet-semilattice with a greatest element $\top$, that is, a partial order such that for all $x, y \in P$ there is a *meet* $x \wedge y$ (an element such that $z \leq x$ and $z \leq y$ implies $z \leq x \wedge y$), and $x \leq \top$ for all x. Then P determines a monoidal category whose 1-cells are the elements of P, and there is a unique 2-cell $(x_1, \ldots, x_n) \Rightarrow (y_1, \ldots, y_m)$ if and only if $x_1 \wedge \ldots \wedge x_n \leq y_1 \wedge \ldots \wedge y_m$:

- the identity 1-cell is $\top$;

- composition and identity 2-cells are the only possible ones;

- for each pair (x, y) of elements, their composite is $x \wedge y$, and the compositor is the unique 2-cell $(x, y) \Rightarrow (x \wedge y)$;

- the unitors and their inverses are given by $\top \wedge x = x$ and $x \wedge \top = x$, which hold for all $x \in P$.

We will meet many more examples of monoidal categories in the course. Now, instead, we focus on the "canonical" example of a bicategory: the bicategory $\underline{\mathbf{Cat}}$ of categories, functors, and *natural transformations*. Recall that $\vec{I}$ is the free category on the graph

$$\overset{0}{\bullet} \xrightarrow{\quad a \quad} \overset{1}{\bullet} \quad .$$

Definition 14. Let X and Y be two categories, $F : X \to Y$ and $G : X \to Y$ two parallel functors. A *natural transformation* $F \Rightarrow G$ is a functor $H : X \times \vec{I} \to Y$ such that $H(-, 0) = F$ and $H(-, 1) = G$.

Remark 15. If you know a little algebraic topology, you can recognise that a natural transformation of functors is defined in the same way as a homotopy of continuous maps, with the category $\vec{I}$ replacing the topological interval. You should keep this intuition in mind: a natural transformation is a way of "continuously moving" from one functor to another.

[Lecture 1, Exercise 37] implies that the product $X \times Y$ of two categories can be given an explicit generators-and-relations presentation, whose generators are the morphisms of the form (f, id_y) and (id_x, g). Therefore, to define a natural transformation $H : X \times \vec{I} \to Y$, it suffices to give:

- $H(f, \mathrm{id}_0)$ and $H(f, \mathrm{id}_1)$ for all $f \in X_1$, but those are already defined to be $F(f)$ and $G(f)$, respectively, and
- $H(\mathrm{id}_x, a) : H(x, 0) \to H(x, 1)$, that is, a morphism $F(x) \to G(x)$ of Y, for each object x of X.

Now, the equation $H(f; f', \mathrm{id}_b) = H(f, \mathrm{id}_b)$ for $b = 0, 1$ is implied by the functoriality of F and G, and there are no non-trivial composites in $\vec{I}$, so the only remaining condition is that, for all morphisms $f : x \to y$ in X,

$$H(f, \mathrm{id}_0); H(\mathrm{id}_y, a) = H(\mathrm{id}_x, a); H(f, \mathrm{id}_1),$$

that is,

$$F(f); H(\mathrm{id}_y, a) = H(\mathrm{id}_x, a); G(f).$$

Letting $\alpha_x := H(\mathrm{id}_x, a)$ for each $x \in X_0$, we can state the following.

Proposition 16. *Let $F, G : X \to Y$ be two functors of categories. The following are equivalent:*

- *a natural transformation $F \Rightarrow G$;*
- *a family $\{\alpha_x : F(x) \to G(x)\}$ of morphisms of Y, parametrised by the objects of X, such that for all morphisms $f : x \to y$ in X, the following is a commutative diagram in Y:*

$$\begin{tikzcd} & G(x) \arrow[dr, "G(f)"] \arrow[dd, phantom, "\Vert"] & \\ F(x) \arrow[ur, "\alpha_x"] \arrow[dr, "F(f)"'] & & G(y) \\ & F(y) \arrow[ur, "\alpha_y"'] & \end{tikzcd} . \tag{4}$$

The diagram (4) is often called a *naturality square*. The α_x are called the *components* of the natural transformation.

Example 17. Let V be a vector space over a field k. There is a vector space V^*, the *dual* of V, whose vectors are the linear functionals $V \to k$; this assignment extends to a contravariant functor $(-)^* : \mathbf{Vec}_k{}^{\mathrm{op}} \to \mathbf{Vec}_k$, sending the linear map $a : V \to W$ to the linear map $a^* : W^* \to V^*$ which maps $f : W \to k$ to $a; f : V \to k$.

Iterating this functor twice gives an *endofunctor* $(-)^{**} : \mathbf{Vec}_k \to \mathbf{Vec}_k$, the *double dual.* Now, for each V, there is a linear map $\lambda_V : V \to V^{**}$ which sends a vector $v \in V$ to the linear functional $\lambda_V(v) : V^* \to k$ so defined: given a functional $f : V \to k$,

$$\lambda_V(v)(f) := f(v).$$

You can check that the family $\{\lambda_V : V \to V^{**}\}$ determines a natural transformation $\lambda : \mathrm{id}_{\mathbf{Vec}_k} \Rightarrow (-)^{**}$.

Example 18. Let X be a category. For all morphisms $f : x \to y$ and objects z of X, there is a function $f_* : \mathrm{Hom}_X(z, x) \to \mathrm{Hom}_X(z, y)$ of sets, which sends the morphism $g : z \to x$ to $g; f : z \to y$. Letting z vary in X, this determines a natural transformation

$$f_* : \mathrm{Hom}_X(-, x) \Rightarrow \mathrm{Hom}_X(-, y)$$

of functors $X^{\mathrm{op}} \to \mathbf{Set}$.

Exercise 19. What is a natural transformation between two functors $1 \to X$? And between two functors $\vec{I} \to X$?

Exercise 20. The opposite category construction defines an endofunctor $(-)^{\mathrm{op}} : \mathbf{Cat} \to \mathbf{Cat}$. If we see groups as one-object groupoids, this restricts to an endofunctor $(-)^{\mathrm{op}} : \mathbf{Grp} \to \mathbf{Grp}$ on the category of groups. Construct a natural transformation

$$\mathrm{id}_{\mathbf{Grp}} \Rightarrow (-)^{\mathrm{op}}.$$

Construction 21. There is a bicategory $\underline{\mathbf{Cat}}$ whose 0-cells are (small) categories, 1-cells are functors, and 2-cells

$$\begin{tikzcd}
& \bullet \arrow[r, dotted] & \bullet \arrow[dr, "G_m"] & \\
\bullet \arrow[ur, "G_1"] \arrow[dr, "F_1"'] \arrow[rrr, phantom, "\Uparrow \alpha"] & & & \bullet \\
& \bullet \arrow[r, dotted] & \bullet \arrow[ur, "F_n"'] &
\end{tikzcd}$$

are natural transformations $\alpha : F_1; \ldots ; F_n \Rightarrow G_1; \ldots ; G_m$, for all composable sequences $(F_1, \ldots , F_n)$ and $(G_1, \ldots , G_m)$ of functors.

The identity 1-cell on a category X is the identity functor id_X, and the identity 2-cell on $(F_1, \ldots , F_n)$ is the *identity natural transformation* from $F_1; \ldots ; F_n$ to itself, whose components are all identities.

The composite in **Cat** of the four diagrams

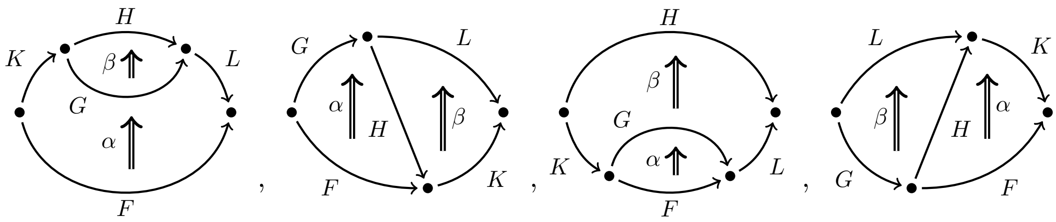

is, respectively,

1. the natural transformation $\alpha; L(\beta_{K-}) : F \Rightarrow K; H; L$ whose components are

$$F(x) \xrightarrow{\alpha_x} LGK(x) \xrightarrow{L(\beta_{K(x)})} LHK(x) \quad ,$$

2. the natural transformation $K(\alpha); \beta_{G-} : F; K \Rightarrow G; L$ whose components are

$$KF(x) \xrightarrow{K(\alpha_x)} KHG(x) \xrightarrow{\beta_{G(x)}} LG(x) \quad ,$$

3. the natural transformation $L(\alpha_{K-}); \beta : K; F; L \Rightarrow H$ whose components are

$$LFK(x) \xrightarrow{L(\alpha_{K(x)})} LGK(x) \xrightarrow{\beta_x} H(x) \quad ,$$

4. the natural transformation $\alpha_{G-}; K(\beta) : G; F \Rightarrow L; K$ whose components are

$$FG(x) \xrightarrow{\alpha_{G(x)}} KHG(x) \xrightarrow{K(\beta_x)} KL(x) \quad .$$

Any other diagrams falls into one of the four situations after composing 1-cells where possible, or introducing identity 1-cells.

Let us prove that the third one is, indeed, a natural transformation; the proof for the other ones is analogous. Because functors preserve all commutative diagrams, each naturality square

$$\begin{tikzcd}
& GK(x) \arrow[dr, "GK(f)"] \arrow[dd, phantom, "\|"] & \\
FK(x) \arrow[ur, "\alpha_{K(x)}"] \arrow[dr, "FK(f)"'] & & GK(y) \\
& FK(y) \arrow[ur, "\alpha_{K(y)}"'] &
\end{tikzcd}$$

is mapped by L to a naturality square

$$\begin{tikzcd}
& LGK(x) \arrow[dr, "LGK(f)"] \arrow[dd, phantom, "\|"] & \\
LFK(x) \arrow[ur, "L(\alpha_{K(x)})"] \arrow[dr, "LFK(f)"'] & & LGK(y) \\
& LFK(y) \arrow[ur, "L(\alpha_{K(y)})"'] &
\end{tikzcd}$$

and because commutative diagrams compose, from each naturality square of β we obtain a commutative diagram

$$\begin{tikzcd}
& & H(x) \arrow[dr, "H(f)"] \arrow[dd, phantom, "\|"] & \\
& LGK(x) \arrow[ur, "\beta_x"] \arrow[dr, "LGK(f)"] \arrow[dd, phantom, "\|"] & & H(y) \\
LFK(x) \arrow[ur, "L(\alpha_{K(x)})"] \arrow[dr, "LFK(f)"'] & & LGK(y) \arrow[ur, "\beta_y"'] & \\
& LFK(y) \arrow[ur, "L(\alpha_{K(y)})"'] & &
\end{tikzcd}$$

and we are done.

Finally, the compositor 2-cells $(F, G) \Rightarrow (F; G)$ and the unitor 2-cells $(\mathrm{id}_X, F) \Rightarrow (F)$ and $(F, \mathrm{id}_Y) \Rightarrow (F)$ are both identity natural transformations.

Exercise 22. Prove that the following are equivalent:

- a natural transformation $F \Rightarrow G$ which is invertible as a 2-cell of $\underline{\mathbf{Cat}}$;
- a natural transformation $F \Rightarrow G$ whose components $F(x) \to G(x)$ are all isomorphisms.

In a certain sense, except for those parts in which one has to deal with categories of different set-theoretic sizes, all of category theory happens inside the bicategory $\underline{\mathbf{Cat}}$. It is not surprising, then, that many of the concepts and constructions of category theory are instances of concepts and constructions that can be formulated in other bicategories.

The fact that bicategories, in general, admit a 2-dimensional calculus of pasting diagrams (later we will also introduce its dual, the calculus of *string diagrams*) is what makes diagrammatic reasoning powerful in category theory.

Construction 23. While a category has only one dual, its opposite category, for a bicategory X there are three different possibilities:

1. the bicategory X^{op} whose 1-cells are reversed but 2-cells are not, that is, for all 1-cells $f : x \to y$ of X there is a 1-cell $f^{\mathrm{op}} : y \to x$ in X^{op}, and for all 2-cells $\alpha : (f_1, \ldots, f_n) \Rightarrow (g_1, \ldots, g_m)$ of X, there is a 2-cell $\alpha^{\mathrm{op}} : ({f_n}^{\mathrm{op}}, \ldots, {f_1}^{\mathrm{op}}) \Rightarrow ({g_m}^{\mathrm{op}}, \ldots, {g_1}^{\mathrm{op}})$ in X^{op};

2. the bicategory X^{co} whose 2-cells are reversed but 1-cells are not, that is, for all 2-cells $\alpha : (f_1, \ldots, f_n) \Rightarrow (g_1, \ldots, g_m)$ of X, there is a 2-cell $\alpha^{\mathrm{co}} : (g_1, \ldots, g_m) \Rightarrow (f_1, \ldots, f_n)$ in X^{co};

3. the bicategory X° whose 1-cells and 2-cells are both reversed, that is, for all 1-cells $f : x \to y$ of X there is a 1-cell $f^{\circ} : y \to x$ in X°, and for all 2-cells $\alpha : (f_1, \ldots, f_n) \Rightarrow (g_1, \ldots, g_m)$ of X, there is a 2-cell $\alpha^{\circ} : ({g_m}^{\circ}, \ldots, {g_1}^{\circ}) \Rightarrow ({f_n}^{\circ}, \ldots, {f_1}^{\circ})$ in X°.

Of course, $X^{\circ} = (X^{\mathrm{op}})^{\mathrm{co}} = (X^{\mathrm{co}})^{\mathrm{op}}$.

We define a particular notion of morphism between bicategories (arguably, not the most common or useful one), which preserves all the structure.

Definition 24. Let X, Y be two bicategories. A *strict functor* of bicategories is a triple of functions $F_i : X_i \to Y_i$, for $i = 0, 1, 2$, such that

1. $s(F(f)) = F(s(f))$ and $t(F(f)) = F(t(f))$ for all $f \in X_1$, implying that F maps X_1^+ to Y_1^+;

2. $s(F(\alpha)) = F(s(\alpha))$ and $t(F(\alpha)) = F(t(\alpha))$ for all $\alpha \in X_2$;

These two imply that F maps pasting diagrams in X to pasting diagrams in Y.

3. F preserves identity 1-cells, 2-cells, and composites of pasting diagrams;

4. F preserves compositors and unitors, that is, $F(c_{f,g}) = c_{F(f),F(g)}$ and $F(l_f) = l_{F(f)}$ and $F(r_f) = r_{F(f)}$.

Example 25. Recall that the functor $(-)^{\mathrm{op}} : \mathbf{Cat} \to \mathbf{Cat}$ assigns to each functor $F : X \to Y$ a functor $F : X^{\mathrm{op}} \to Y^{\mathrm{op}}$. What about natural transformations $\alpha : F \Rightarrow G$?

In components, α is given by a family of morphisms $\alpha_x : F(x) \to G(x)$ in Y. Through $(-)^{\mathrm{op}}$, each of these is mapped to a morphism $\alpha_x^{\mathrm{op}} : G(x) \to F(x)$ in Y^{op}. You can check that these define a natural transformation from $G : X^{\mathrm{op}} \to Y^{\mathrm{op}}$ to $F : X^{\mathrm{op}} \to Y^{\mathrm{op}}$: that is, $(-)^{\mathrm{op}}$ *reverses* the direction of a natural transformation.

This is compatible with composition, and $(-)^{\mathrm{op}}$ extends to a strict functor

$$(-)^{\mathrm{op}} : \underline{\mathbf{Cat}}^{\mathrm{co}} \to \underline{\mathbf{Cat}}$$

of bicategories, which is “contravariant in the direction of 2-cells”.

This fact is at the heart of the fundamental *duality principle* in category theory: the fact that any true statement about categories, functors, and natural transformations has a *dual* true statement involving *opposite* categories, functors, and *opposite* natural transformations, obtained by applying $(-)^{\mathrm{op}}$ to whatever is involved.

Basically, this means that every theorem in category theory is actually two theorems in one.