2

Adjoints

The slogan of Saunders Mac Lane's book *Categories for the Working Mathematician* is:

> *Adjoint functors arise everywhere.*

We will see the truth of this, meeting examples of adjoint functors from diverse parts of mathematics. To complement the understanding provided by examples, we will approach the theory of adjoints from three different directions, each of which carries its own intuition. Then we will prove that the three approaches are equivalent.

Understanding adjointness gives you a valuable addition to your mathematical toolkit. Most professional pure mathematicians know what categories and functors are, but far fewer know about adjoints. More should: adjoint functors are both common and easy, and knowing about adjoints helps you to spot patterns in the mathematical landscape.

2.1 Definition and examples

Consider a pair of functors in opposite directions, $F\colon \mathscr{A} \to \mathscr{B}$ and $G\colon \mathscr{B} \to \mathscr{A}$. Roughly speaking, F is said to be left adjoint to G if, whenever $A \in \mathscr{A}$ and $B \in \mathscr{B}$, maps $F(A) \to B$ are essentially the same thing as maps $A \to G(B)$.

Definition 2.1.1 Let $\mathscr{A} \underset{G}{\overset{F}{\rightleftarrows}} \mathscr{B}$ be categories and functors. We say that F is **left adjoint** to G, and G is **right adjoint** to F, and write $F \dashv G$, if

$$\mathscr{B}(F(A), B) \cong \mathscr{A}(A, G(B)) \tag{2.1}$$

naturally in $A \in \mathscr{A}$ and $B \in \mathscr{B}$. The meaning of 'naturally' is defined below. An **adjunction** between F and G is a choice of natural isomorphism (2.1).

'Naturally in $A \in \mathscr{A}$ and $B \in \mathscr{B}$' means that there is a specified bijection (2.1) for each $A \in \mathscr{A}$ and $B \in \mathscr{B}$, and that it satisfies a naturality axiom. To state it, we need some notation. Given objects $A \in \mathscr{A}$ and $B \in \mathscr{B}$, the correspondence (2.1) between maps $F(A) \to B$ and $A \to G(B)$ is denoted by a horizontal bar, in both directions:

$$\begin{aligned} \bigl(F(A) \xrightarrow{g} B\bigr) &\quad\mapsto\quad \bigl(A \xrightarrow{\bar{g}} G(B)\bigr), \\ \bigl(F(A) \xrightarrow{\bar{f}} B\bigr) &\quad\leftarrow\!\shortmid\quad \bigl(A \xrightarrow{f} G(B)\bigr). \end{aligned}$$

So $\bar{\bar{f}} = f$ and $\bar{\bar{g}} = g$. We call $\bar{f}$ the **transpose** of f, and similarly for g. The naturality axiom has two parts:

$$\overline{\bigl(F(A) \xrightarrow{g} B \xrightarrow{q} B'\bigr)} \quad = \quad \bigl(A \xrightarrow{\bar{g}} G(B) \xrightarrow{G(q)} G(B')\bigr) \tag{2.2}$$

(that is, $\overline{q \circ g} = G(q) \circ \bar{g}$) for all g and q, and

$$\overline{\bigl(A' \xrightarrow{p} A \xrightarrow{f} G(B)\bigr)} \quad = \quad \bigl(F(A') \xrightarrow{F(p)} F(A) \xrightarrow{\bar{f}} B\bigr) \tag{2.3}$$

for all p and f. It makes no difference whether we put the long bar over the left or the right of these equations, since bar is self-inverse.

Remarks 2.1.2 (a) The naturality axiom might seem ad hoc, but we will see in Chapter 4 that it simply says that two particular functors are naturally isomorphic. In this section, we ignore the naturality axiom altogether, trusting that it embodies our usual intuitive idea of naturality: something defined without making any arbitrary choices.

(b) The naturality axiom implies that from each array of maps

$$A_0 \to \cdots \to A_n, \quad F(A_n) \to B_0, \quad B_0 \to \cdots \to B_m,$$

it is possible to construct exactly one map

$$A_0 \to G(B_m).$$

Compare the comments on the definitions of category, functor and natural transformation (Remarks 1.1.2(b), 1.2.2(a), and 1.3.2(a)).

(c) Not only do adjoint functors arise everywhere; better, whenever you see a pair of functors $\mathscr{A} \rightleftarrows \mathscr{B}$, there is an excellent chance that they are adjoint (one way round or the other).

For example, suppose you get talking to a mathematician who tells you that her work involves Lie algebras and associative algebras. You try to object that you don't know what either of those things is, but she carries on talking anyway, explaining that there's a way of turning any Lie algebra into an associative algebra, and also a way of turning any associative algebra into a Lie algebra. At this point, even without knowing what she's talking about, you should bet her that one process is adjoint to the other. This almost always works.

(d) A given functor G may or may not have a left adjoint, but if it does, it is unique up to isomorphism, so we may speak of '*the* left adjoint of G'. The same goes for right adjoints. We prove this later (Example 4.3.13).

You might ask 'what do we gain from knowing that two functors are adjoint?' The uniqueness is a crucial part of the answer. Let us return to the example of (c). It would take you only a few minutes to learn what Lie algebras are, what associative algebras are, and what the standard functor G is that turns an associative algebra into a Lie algebra. What about the functor F in the opposite direction? The description of F that you will find in most algebra books (under 'universal enveloping algebra') takes much longer to understand. However, you can bypass that process completely, just by knowing that F is the left adjoint of G. Since G can have only *one* left adjoint, this characterizes F completely. In a sense, it tells you all you need to know.

Examples 2.1.3 (Algebra: free $\dashv$ forgetful) Forgetful functors between categories of algebraic structures usually have left adjoints. For instance:

(a) Let k be a field. There is an adjunction

$$\begin{tikzcd} \mathbf{Vect}_k \arrow[d, shift left, "U"] \\ \mathbf{Set}, \arrow[u, shift left, "F"] \arrow[u, phantom, "\dashv"] \end{tikzcd}$$

where U is the forgetful functor of Example 1.2.3(b) and F is the free functor of Example 1.2.4(c). Adjointness says that given a set S and a vector space V, a linear map $F(S) \to V$ is essentially the same thing as a function $S \to U(V)$.

We saw this in Example 0.4, but let us now check it in detail.

Fix a set S and a vector space V. Given a linear map $g: F(S) \to V$, we may define a map of sets $\bar{g}: S \to U(V)$ by $\bar{g}(s) = g(s)$ for all $s \in S$. This gives a function

$$\begin{array}{ccc} \mathbf{Vect}_k(F(S), V) & \to & \mathbf{Set}(S, U(V)) \\ g & \mapsto & \bar{g}. \end{array}$$

In the other direction, given a map of sets $f: S \to U(V)$, we may define a linear map $\bar{f}: F(S) \to V$ by $\bar{f}(\sum_{s\in S} \lambda_s s) = \sum_{s\in S} \lambda_s f(s)$ for all formal linear combinations $\sum \lambda_s s \in F(S)$. This gives a function

$$\begin{array}{ccc} \mathbf{Set}(S, U(V)) & \to & \mathbf{Vect}_k(F(S), V) \\ f & \mapsto & \bar{f}. \end{array}$$

These two functions 'bar' are mutually inverse: for any linear map $g\colon F(S) \to V$, we have

$$\bar{\bar{g}}\left(\sum_{s\in S} \lambda_s s\right) = \sum_{s\in S} \lambda_s \bar{g}(s) = \sum_{s\in S} \lambda_s g(s) = g\left(\sum_{s\in S} \lambda_s s\right)$$

for all $\sum \lambda_s s \in F(S)$, so $\bar{\bar{g}} = g$, and for any map of sets $f\colon S \to U(V)$, we have

$$\bar{\bar{f}}(s) = \bar{f}(s) = f(s)$$

for all $s \in S$, so $\bar{\bar{f}} = f$. We therefore have a canonical bijection between $\mathbf{Vect}_k(F(S), V)$ and $\mathbf{Set}(S, U(V))$ for each $S \in \mathbf{Set}$ and $V \in \mathbf{Vect}_k$, as required.

Here we have been careful to distinguish between the vector space V and its underlying set $U(V)$. Very often, though, in category theory as in mathematics at large, the symbol for a forgetful functor is omitted. In this example, that would mean dropping the U and leaving the reader to figure out whether each occurrence of V is intended to denote the vector space itself or its underlying set. We will soon start using such notational shortcuts ourselves.

(b) In the same way, there is an adjunction

$$\begin{tikzcd} \mathbf{Grp} \arrow[d, shift left, "U"] \arrow[d, phantom, "\dashv"] \\ \mathbf{Set} \arrow[u, shift left, "F"] \end{tikzcd}$$

where F and U are the free and forgetful functors of Examples 1.2.3(a) and 1.2.4(a).

The free group functor is tricky to construct explicitly. In Chapter 6, we will prove a result (the general adjoint functor theorem) guaranteeing that U and many functors like it all have left adjoints. To some extent, this removes the need to construct F explicitly, as observed in Remark 2.1.2(d). The point can be overstated: for a group theorist, the more descriptions of free groups that are available, the better. Explicit constructions really can be useful. But it is an important general principle that forgetful functors of this type always have left adjoints.

(c) There is an adjunction

$$\begin{tikzcd} \mathbf{Ab} \arrow[d, shift left, "U"] \\ \mathbf{Grp} \arrow[u, shift left, "F"] \arrow[u, phantom, "\dashv"] \end{tikzcd}$$

where U is the inclusion functor of Example 1.2.3(d). If G is a group then $F(G)$ is the **abelianization** G_{ab} of G. This is an abelian quotient group of G, with the property that every map from G to an abelian group factorizes uniquely through G_{ab}:

$$\begin{tikzcd} G \arrow[r, "\eta"] \arrow[dr, "\forall\phi"'] & G_{\mathrm{ab}} \arrow[d, dotted, "\exists!\bar{\phi}"] \\ & \forall A. \end{tikzcd}$$

Here η is the natural map from G to its quotient G_{ab}, and A is any abelian group. (We have adopted the abuse of notation advertised in example (a), omitting the symbol U at several places in this diagram.) The bijection

$$\mathbf{Ab}(G_{\mathrm{ab}}, A) \cong \mathbf{Grp}(G, U(A))$$

is given in the left-to-right direction by $\psi \mapsto \psi \circ \eta$, and in the right-to-left direction by $\phi \mapsto \bar{\phi}$.

(To construct G_{ab}, let G' be the smallest normal subgroup of G containing $xyx^{-1}y^{-1}$ for all $x, y \in G$, and put $G_{\mathrm{ab}} = G/G'$. The kernel of any homomorphism from G to an abelian group contains G', and the universal property follows.)

(d) There are adjunctions

$$\begin{tikzcd} \mathbf{Grp} \arrow[d, "U"] \\ \mathbf{Mon} \arrow[u, bend left, "F"] \arrow[u, bend right, "R"'] \end{tikzcd} \qquad F \dashv U \dashv R$$

between the categories of groups and monoids. The middle functor U is inclusion. The left adjoint F is, again, tricky to describe explicitly. Informally, $F(M)$ is obtained from M by throwing in an inverse to every element. (For example, if M is the additive monoid of natural numbers then $F(M)$ is the group of integers.) Again, the general adjoint functor theorem (Theorem 6.3.10) guarantees the existence of this adjoint.

This example is unusual in that forgetful functors do not usually have *right* adjoints. Here, given a monoid M, the group $R(M)$ is the submonoid of M consisting of all the invertible elements.

The category **Grp** is both a **reflective** and a **coreflective** subcategory of **Mon**. This means, by definition, that the inclusion functor $\mathbf{Grp} \hookrightarrow \mathbf{Mon}$ has both a left and a right adjoint. The previous example tells us that **Ab** is a reflective subcategory of **Grp**.

(e) Let **Field** be the category of fields, with ring homomorphisms as the maps. The forgetful functor $\mathbf{Field} \to \mathbf{Set}$ does *not* have a left adjoint. (For a proof, see Example 6.3.5.) The theory of fields is unlike the theories of groups, rings, and so on, because the operation $x \mapsto x^{-1}$ is not defined for *all* x (only for $x \neq 0$).

Remark 2.1.4 At several points in this book, we make contact with the idea of an **algebraic theory**. You already know several examples: the theory of groups is an algebraic theory, as are the theory of rings, the theory of vector spaces over $\mathbb{R}$, the theory of vector spaces over $\mathbb{C}$, the theory of monoids, and (rather trivially) the theory of sets. After reading the description below, you might conclude that the word 'theory' is overly grand, and that 'definition' would be more appropriate. Nevertheless, this is the established usage.

We will not need to define 'algebraic theory' formally, but it will be important to have the general idea. Let us begin by considering the theory of groups.

A group can be defined as a set X equipped with a function $\cdot\colon X \times X \to X$ (multiplication), another function $(\ )^{-1}\colon X \to X$ (inverse), and an element $e \in X$ (the identity), satisfying a familiar list of equations. More systematically, the three pieces of structure on X can be seen as maps of sets

$$\cdot\colon X^2 \to X, \qquad (\ )^{-1}\colon X^1 \to X, \qquad e\colon X^0 \to X,$$

where in the last case, X^0 is the one-element set 1 and we are using the observation that a map $1 \to X$ of sets is essentially the same thing as an element of X.

(You may be more familiar with a definition of group in which only the multiplication and perhaps the identity are specified as pieces of *structure*, with the existence of inverses required as a *property*. In that approach, the definition is swiftly followed by a lemma on uniqueness of inverses, guaranteeing that it makes sense to speak of *the* inverse of an element. The two approaches are equivalent, but for many purposes, it is better to frame the definition in the way described in the previous paragraph.)

An algebraic theory consists of two things: first, a collection of operations, each with a specified arity (number of inputs), and second, a collection of equations. For example, the theory of groups has one operation of arity 2, one of arity 1, and one of arity 0. An **algebra** or **model** for an algebraic theory consists of a set X together with a specified map $X^n \to X$ for each operation of arity n, such that the equations hold everywhere. For example, an algebra for the theory of groups is exactly a group.

A more subtle example is the theory of vector spaces over $\mathbb{R}$. This is an algebraic theory with, among other things, an infinite number of operations of arity 1: for each $\lambda \in \mathbb{R}$, we have the operation $\lambda \cdot -: X \to X$ of scalar multiplication by λ (for any vector space X). There is nothing special about the field $\mathbb{R}$ here; the only point is that it was chosen in advance. The theory of vector spaces over $\mathbb{R}$ is different from the theory of vector spaces over $\mathbb{C}$, because they have different operations of arity 1.

In a nutshell, the main property of algebras for an algebraic theory is that the operations are defined everywhere on the set, and the equations hold everywhere too. For example, *every* element of a group has a specified inverse, and *every* element x satisfies the equation $x \cdot x^{-1} = 1$. This is why the theories of groups, rings, and so on, are algebraic theories, but the theory of fields is not.

Example 2.1.5 There are adjunctions

$$\begin{tikzcd}
\mathbf{Top} \arrow[d, "U"] \\
\mathbf{Set} \arrow[u, bend left=50, "D"] \arrow[u, bend right=50, "I"']
\end{tikzcd} \qquad D \dashv U \dashv I$$

where U sends a space to its set of points, D equips a set with the discrete topology, and I equips a set with the indiscrete topology.

Example 2.1.6 Given sets A and B, we can form their (cartesian) product $A \times B$. We can also form the set B^A of functions from A to B. This is the same as the set $\mathbf{Set}(A, B)$, but we tend to use the notation B^A when we want to emphasize that it is an object of the same category as A and B.

Now fix a set B. Taking the product with B defines a functor

$$\begin{array}{rccc} - \times B\colon & \mathbf{Set} & \to & \mathbf{Set} \\ & A & \mapsto & A \times B. \end{array}$$

(Here we are using the blank notation introduced in Example 1.2.12.) There is also a functor

$$\begin{array}{rccc} (-)^B\colon & \mathbf{Set} & \to & \mathbf{Set} \\ & C & \mapsto & C^B. \end{array}$$

Moreover, there is a canonical bijection

$$\mathbf{Set}(A \times B, C) \cong \mathbf{Set}(A, C^B)$$

for any sets A and C. It is defined by simply changing the punctuation: given a

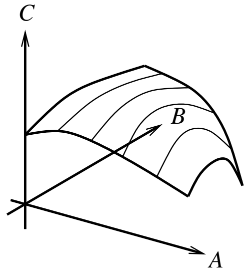

Figure 2.1 In **Set**, a map $A \times B \to C$ can be seen as a way of assigning to each element of A a map $B \to C$.

map $g\colon A \times B \to C$, define $\bar{g}\colon A \to C^B$ by

$$(\bar{g}(a))(b) = g(a, b)$$

($a \in A$, $b \in B$), and in the other direction, given $f\colon A \to C^B$, define $\bar{f}\colon A \times B \to C$ by

$$\bar{f}(a, b) = (f(a))(b)$$

($a \in A$, $b \in B$). Figure 2.1 shows an example with $A = B = C = \mathbb{R}$. By slicing up the surface as shown, a map $\mathbb{R}^2 \to \mathbb{R}$ can be seen as a map from $\mathbb{R}$ to $\{\text{maps } \mathbb{R} \to \mathbb{R}\}$.

Putting all this together, we obtain an adjunction

$$\begin{tikzcd} \mathbf{Set} \arrow[d, shift left, "(-)^B"] \\ \mathbf{Set} \arrow[u, shift left, "- \times B"] \arrow[u, phantom, "\dashv"] \end{tikzcd}$$

for every set B.

Definition 2.1.7 Let $\mathscr{A}$ be a category. An object $I \in \mathscr{A}$ is **initial** if for every $A \in \mathscr{A}$, there is exactly one map $I \to A$. An object $T \in \mathscr{A}$ is **terminal** if for every $A \in \mathscr{A}$, there is exactly one map $A \to T$.

For example, the empty set is initial in **Set**, the trivial group is initial in **Grp**, and $\mathbb{Z}$ is initial in **Ring** (Example 0.2). The one-element set is terminal in **Set**, the trivial group is terminal (as well as initial) in **Grp**, and the trivial (one-element) ring is terminal in **Ring**. The terminal object of **CAT** is the category **1** containing just one object and one map (necessarily the identity on that object).

A category need not have an initial object, but if it does have one, it is unique up to isomorphism. Indeed, it is unique up to *unique* isomorphism, as follows.

Lemma 2.1.8 *Let I and I' be initial objects of a category. Then there is a unique isomorphism $I \to I'$. In particular, $I \cong I'$.*

Proof Since I is initial, there is a unique map $f\colon I \to I'$. Since I' is initial, there is a unique map $f'\colon I' \to I$. Now $f' \circ f$ and 1_I are both maps $I \to I$, and I is initial, so $f' \circ f = 1_I$. Similarly, $f \circ f' = 1_{I'}$. Hence f is an isomorphism, as required. □

Example 2.1.9 Initial and terminal objects can be described as adjoints. Let $\mathscr{A}$ be a category. There is precisely one functor $\mathscr{A} \to \mathbf{1}$. Also, a functor $\mathbf{1} \to \mathscr{A}$ is essentially just an object of $\mathscr{A}$ (namely, the object to which the unique object of $\mathbf{1}$ is mapped). Viewing functors $\mathbf{1} \to \mathscr{A}$ as objects of $\mathscr{A}$, a left adjoint to $\mathscr{A} \to \mathbf{1}$ is exactly an initial object of $\mathscr{A}$.

Similarly, a right adjoint to the unique functor $\mathscr{A} \to \mathbf{1}$ is exactly a terminal object of $\mathscr{A}$.

Remark 2.1.10 In the language introduced in Remark 1.1.10, the concept of terminal object is dual to the concept of initial object. (More generally, the concepts of left and right adjoint are dual to one another.) Since any two initial objects of a category are uniquely isomorphic, the principle of duality implies that the same is true of terminal objects.

Remark 2.1.11 Adjunctions can be composed. Take adjunctions

$$\begin{tikzcd} \mathscr{A} \arrow[r, shift left=1.5, "F"] \arrow[r, phantom, "\perp"] & \mathscr{A}' \arrow[l, shift left=1.5, "G"] \arrow[r, shift left=1.5, "F'"] \arrow[r, phantom, "\perp"] & \mathscr{A}'' \arrow[l, shift left=1.5, "G'"] \end{tikzcd}$$

where the $\perp$ symbol is a rotated $\dashv$ (thus, $F \dashv G$ and $F' \dashv G'$). Then we obtain an adjunction

$$\begin{tikzcd} \mathscr{A} \arrow[r, shift left=1.5, "F' \circ F"] \arrow[r, phantom, "\perp"] & \mathscr{A}'' \arrow[l, shift left=1.5, "G \circ G'"] \end{tikzcd},$$

since for $A \in \mathscr{A}$ and $A'' \in \mathscr{A}''$,

$$\mathscr{A}''(F'(F(A)), A'') \cong \mathscr{A}'(F(A), G'(A'')) \cong \mathscr{A}(A, G(G'(A'')))$$

naturally in A and A''.

Exercises

2.1.12 Find three examples of adjoint functors not mentioned above. Do the same for initial and terminal objects.

2.1.13 What can be said about adjunctions between discrete categories?

2.1.14 Show that the naturality equations (2.2) and (2.3) can equivalently be replaced by the single equation

$$\overline{\left(A' \xrightarrow{p} A \xrightarrow{f} G(B) \xrightarrow{G(q)} G(B')\right)} \quad = \quad \left(F(A') \xrightarrow{F(p)} F(A) \xrightarrow{\bar{f}} B \xrightarrow{q} B'\right)$$

for all p, f and q.

2.1.15 Show that left adjoints preserve initial objects: that is, if $\mathscr{A} \underset{G}{\overset{F}{\rightleftarrows}} \mathscr{B}$ with $F \dashv G$ and I is an initial object of $\mathscr{A}$, then $F(I)$ is an initial object of $\mathscr{B}$. Dually, show that right adjoints preserve terminal objects.

(In Section 6.3, we will see this as part of a bigger picture: right adjoints preserve limits and left adjoints preserve colimits.)

2.1.16 Let G be a group.

(a) What interesting functors are there (in either direction) between **Set** and the category $[G, \mathbf{Set}]$ of left G-sets? Which of those functors are adjoint to which?

(b) Similarly, what interesting functors are there between $\mathbf{Vect}_k$ and the category $[G, \mathbf{Vect}_k]$ of k-linear representations of G, and what adjunctions are there between those functors?

2.1.17 Fix a topological space X, and write $\mathscr{O}(X)$ for the poset of open subsets of X, ordered by inclusion. Let

$$\Delta\colon \mathbf{Set} \to [\mathscr{O}(X)^{\mathrm{op}}, \mathbf{Set}]$$

be the functor assigning to a set A the presheaf ΔA with constant value A. Exhibit a chain of adjoint functors

$$\Lambda \dashv \Pi \dashv \Delta \dashv \Gamma \dashv \nabla.$$

2.2 Adjunctions via units and counits

In the previous section, we met the definition of adjunction. In this section and the next, we meet two ways of rephrasing the definition. The one in this section is most useful for theoretical purposes, while the one in the next fits well with many examples.

To start building the theory of adjoint functors, we have to take seriously the naturality requirement (equations (2.2) and (2.3)), which has so far been