

Abelianization Volley #3 - Adjoint Functors

Tucker Q · 2 Jun 2024

Welcome back to the series! This time we'll be discussing a concept I've repeatedly alluded to, but deliberately danced around: that of adjunctions / adjoint functors. Adjunctions are one of the most common and useful constructions in category theory. In the words of Saunders Mac Lane, "Adjoint functors arise everywhere" - and we will see several examples in the course of this series. Before we dive in to the definition and start proving theorems, I want to provide a high-level overview.

Suppose we have categories \mathbf{C} and \mathbf{D} . An adjunction from \mathbf{D} to \mathbf{C} consists of a pair of functors $F : \mathbf{D} \rightarrow \mathbf{C}$ and $G : \mathbf{C} \rightarrow \mathbf{D}$, which are related in a way we will make explicit later. The functor F is called the left adjoint, and G the right adjoint - this relationship may be written $F \dashv G$. An adjunction is often described as a weaker form of equivalence of categories. This is mainly because the functors which constitute an equivalence of categories each preserve both limits and colimits, but we will see later that for adjunctions, F preserves colimits and G preserves limits. However, I think this characterization can be a bit misleading: equivalence is usually thought of as a relationship between two categories, while adjunction is better conceptualized as a relationship between a pair of functors. The functors themselves make up an "adjoint pair", and the relationship between the ambient categories is usually less central.

Definition 1. Let \mathbf{C} and \mathbf{D} be categories. An adjunction from \mathbf{D} to \mathbf{C} is an ordered triple $\langle F, G, \varphi \rangle$, where $F : \mathbf{D} \rightarrow \mathbf{C}$ and $G : \mathbf{C} \rightarrow \mathbf{D}$ are functors, and $\varphi : \text{Hom}_{\mathbf{C}}(F(-), -) \rightarrow \text{Hom}_{\mathbf{D}}(-, G(-))$ is an isomorphism which is natural in each of its arguments.

It is important to remember that hom-functors are contravariant in the first argument and covariant in the second - this means that the diagrams expressing the naturality of φ in each argument look different. Let $f \in \text{Hom}_{\mathbf{D}}(X, X')$ and $g \in \text{Hom}_{\mathbf{C}}(A, A')$ be arrows; then for φ to be natural in each argument, the following diagrams must commute. The first diagram expresses naturality in the first argument, and the second diagram expresses naturality in the second.

$$\begin{array}{ccc}
\mathrm{Hom}_{\mathbf{C}}(F(X), A) & \xrightarrow{\varphi_{X,A}} & \mathrm{Hom}_{\mathbf{D}}(X, G(A)) \\
\mathrm{Hom}_{\mathbf{C}}(F(f), A) \uparrow & & \uparrow \mathrm{Hom}_{\mathbf{D}}(f, G(A)) \\
\mathrm{Hom}_{\mathbf{C}}(F(X'), A) & \xrightarrow{\varphi_{X',A}} & \mathrm{Hom}_{\mathbf{D}}(X', G(A))
\end{array} \tag{1}$$

$$\begin{array}{ccc}
\mathrm{Hom}_{\mathbf{C}}(F(X), A) & \xrightarrow{\varphi_{X,A}} & \mathrm{Hom}_{\mathbf{D}}(X, G(A)) \\
\mathrm{Hom}_{\mathbf{C}}(F(X), g) \downarrow & & \downarrow \mathrm{Hom}_{\mathbf{D}}(X, G(g)) \\
\mathrm{Hom}_{\mathbf{C}}(F(X), A') & \xrightarrow{\varphi_{X,A'}} & \mathrm{Hom}_{\mathbf{D}}(X, G(A'))
\end{array}$$

With this in mind, we are ready to see an example extended from some previous discussion. In the first post in this series, we defined the diagonal functor $\Delta : \mathbf{C} \rightarrow \mathbf{C} \times \mathbf{C}$ which maps objects and arrows to pairs thereof. Now suppose that \mathbf{C} is a category which has all binary coproducts. This allows us to define the coproduct as a functor as well:

Definition 2. The coproduct determines a functor $\amalg : \mathbf{C} \times \mathbf{C} \rightarrow \mathbf{C}$ such that for (A, B) in $\mathbf{C} \times \mathbf{C}$, $\amalg(A, B) = A \amalg B$ and for an arrow $(f, g) : (A, B) \rightarrow (A', B')$ in $\mathbf{C} \times \mathbf{C}$, $\amalg(f, g)$ is the unique arrow such that the following diagram commutes, guaranteed by the coproduct property of $A \amalg B$

$$\begin{array}{ccccc}
A & \xrightarrow{\iota_A} & A \amalg B & \xleftarrow{\iota_B} & B \\
f \downarrow & & \amalg(f, g) \downarrow & & \downarrow g \\
A' & \xrightarrow{\iota_{A'}} & A' \amalg B' & \xleftarrow{\iota_{B'}} & B'
\end{array} \tag{2}$$

It is relatively simple to verify that \amalg is indeed a functor. And recall from the end of entry #1 in this series that since $(A \amalg B, (\iota_A, \iota_B))$ is a universal arrow from (A, B) to Δ , we have a natural isomorphism, say φ , from $\mathrm{Hom}_{\mathbf{C}}(\amalg(A, B), -)$ to $\mathrm{Hom}_{\mathbf{C} \times \mathbf{C}}((A, B), \Delta(-))$. For an arrow $f : A \amalg B \rightarrow X$ in \mathbf{C} , this natural isomorphism is determined by $\varphi(f) = \Delta(f) \circ (\iota_A, \iota_B) = (f \circ \iota_A, f \circ \iota_B)$. From here, we would like to show that the functors \amalg and Δ form an adjoint pair $\amalg \dashv \Delta$. We are already most of the way there - all that remains to be shown is that the set of isomorphisms we have denoted by φ also constitute a natural isomorphism in the first argument. To do so, let $(f, g) : (A, B) \rightarrow (A', B')$ be an arrow in $\mathbf{C} \times \mathbf{C}$ and X be a fixed object in \mathbf{C} . We must show that the following diagram commutes:

$$\begin{array}{ccc}
\mathrm{Hom}_{\mathbf{C}}(\amalg(A, B), X) & \xrightarrow{\varphi_{(A,B),X}} & \mathrm{Hom}_{\mathbf{C} \times \mathbf{C}}((A, B), \Delta(X)) \\
\mathrm{Hom}_{\mathbf{C}}(\amalg(f, g), X) \uparrow & & \uparrow \mathrm{Hom}_{\mathbf{C} \times \mathbf{C}}((f, g), \Delta(X)) \\
\mathrm{Hom}_{\mathbf{C}}(\amalg(A', B'), X) & \xrightarrow{\varphi_{(A',B'),X}} & \mathrm{Hom}_{\mathbf{C} \times \mathbf{C}}((A', B'), \Delta(X))
\end{array} \tag{3}$$

To do this, let $h \in \mathrm{Hom}_{\mathbf{C}}(\amalg(A', B'), X)$; we must show that $\varphi(h \circ \amalg(f, g)) = \varphi(h) \circ (f, g)$. Expanding out the definition of φ in each case, this is true if and only if

$$(h \circ \Pi(f, g) \circ \iota_A, h \circ \Pi(f, g) \circ \iota_B) = (h \circ \iota_{A'}, h \circ \iota_{B'}) \circ (f, g) \quad (4)$$

By the definition of $\Pi(f, g)$, we have $\Pi(f, g) \circ \iota_A = \iota_{A'} \circ f$ and $\Pi(f, g) \circ \iota_B = \iota_{B'} \circ g$, so this is clearly true. Therefore the functors Π and Δ do in fact form an adjunction together with the isomorphisms φ . But it turns out that we can make this result even more general: every time we have a universal arrow to a functor from each object of its codomain, that functor will have a left adjoint.

Theorem 1. *Let $G : \mathbf{C} \rightarrow \mathbf{D}$ be a functor, and suppose that for each $X \in \text{ob}(\mathbf{D})$, we have an object X' in \mathbf{C} and an arrow $\eta_X : X \rightarrow G(X')$ such that the pair (X', η_X) is a universal arrow from X to G . Then this set of universal arrows determines a functor $F : \mathbf{D} \rightarrow \mathbf{C}$ and a natural isomorphism $\varphi : \text{Hom}_{\mathbf{C}}(F(-), -) \rightarrow \text{Hom}_{\mathbf{D}}(-, G(-))$ which make up an adjunction $F \dashv G$. Further, the universal arrows η_X are the components of a natural transformation $\eta : 1_{\mathbf{D}} \rightarrow G \circ F$.*

Proof. First, we must construct the functor $F : \mathbf{D} \rightarrow \mathbf{C}$. To define F on objects, we will make the obvious choice $F(X) = X'$. To define F on arrows, let $f : X \rightarrow Y$ be an arrow in \mathbf{D} . Then since $\eta_Y \circ f$ is an arrow from X to $G(Y')$, the universal arrow (X', η_X) gives us a unique arrow $f' : X' \rightarrow Y'$ such that the following diagram commutes:

$$\begin{array}{ccc} X' & X \xrightarrow{\eta_X} & G(X') \\ \downarrow f' & f \downarrow & \downarrow G(f') \\ Y' & Y \xrightarrow{\eta_Y} & G(Y') \end{array} \quad (5)$$

In this case, we will let $F(f) = f'$. Next, we must demonstrate that F is in fact a functor, i.e. that it preserves identities and composition. Seeing that F preserves identities is rather simple - we know that $G(1_{X'}) \circ \eta_X = \eta_X = \eta_X \circ 1_X$. As $F(1_X)$ is defined as the unique arrow from X' to X' with this property, we have $F(1_X) = 1_{X'} = 1_{F(X)}$. To see that F preserves composition of arrows, let $f : X \rightarrow Y$ and $g : Y \rightarrow Z$ be arrows in \mathbf{D} . Then $F(f)$, $F(g)$, and $F(g \circ f)$ are defined as the unique arrows with appropriate domains and codomains such that each of the following diagrams commutes.

$$\begin{array}{ccc} X \xrightarrow{\eta_X} & G(F(X)) & Y \xrightarrow{\eta_Y} & G(F(Y)) \\ f \downarrow & \downarrow G(F(f)) & g \downarrow & \downarrow G(F(g)) \\ Y \xrightarrow{\eta_Y} & G(F(Y)) & Z \xrightarrow{\eta_Z} & G(F(Z)) \end{array} \quad (6)$$

$$\begin{array}{ccc} X \xrightarrow{\eta_X} & G(F(X)) \\ g \circ f \downarrow & \downarrow G(F(g \circ f)) \\ Z \xrightarrow{\eta_Z} & G(F(Z)) \end{array}$$

Using the commutativity of the first two diagrams and the fact that G preserves composition of arrows, we can show that

$$\begin{aligned} G(F(g) \circ F(f)) \circ \eta_X &= G(F(g)) \circ G(F(f)) \circ \eta_X = G(F(g)) \circ \eta_Y \circ f \\ &= \eta_Z \circ g \circ f \end{aligned} \quad (7)$$

But as this is the unique property satisfied by $F(g \circ f)$, we must have $F(g \circ f) = F(g) \circ F(f)$, meaning that F is indeed a functor. At this point, it should also be clear that the arrows η_X make up a natural transformation $\eta : 1_{\mathbf{D}} \rightarrow G \circ F$; this fact is essentially built into our construction of F .

To proceed further, we need to produce the natural isomorphism φ which completes the adjunction. By theorem 2 in the first post in this series, each universal arrow η_X determines a natural isomorphism $\varphi_X : \text{Hom}_{\mathbf{C}}(F(X), -) \rightarrow \text{Hom}_{\mathbf{D}}(X, G(-))$. For an arrow $h : F(X) \rightarrow A$ in \mathbf{C} , the mapping $\varphi_{X,A}$ is defined by $\varphi_{X,A}(h) = G(h) \circ \eta_X$. All that we must do is show that this family of isomorphisms is also natural in the first argument X . To do so, let A be a fixed object in \mathbf{C} , and let $f : X \rightarrow Y$ be an arrow in \mathbf{D} . We must show that the following diagram commutes:

$$\begin{array}{ccc} \text{Hom}_{\mathbf{C}}(F(X), A) & \xrightarrow{\varphi_{X,A}} & \text{Hom}_{\mathbf{D}}(X, G(A)) \\ \text{Hom}_{\mathbf{C}}(F(f), A) \uparrow & & \uparrow \text{Hom}_{\mathbf{D}}(f, G(A)) \\ \text{Hom}_{\mathbf{C}}(F(Y), A) & \xrightarrow{\varphi_{Y,A}} & \text{Hom}_{\mathbf{D}}(Y, G(A)) \end{array} \quad (8)$$

In order to demonstrate this, let $h \in \text{Hom}_{\mathbf{C}}(F(Y), A)$. Along the upper path through the diagram, h is mapped to $\varphi_{X,A}(h \circ F(f))$, which is equal to $G(h \circ F(f)) \circ \eta_X = G(h) \circ G(F(f)) \circ \eta_X$. Along the lower path through the diagram, h is mapped to $\varphi_{Y,A}(h) \circ f$, which is equal to $G(h) \circ \eta_Y \circ f$. But since $\eta_Y \circ f = G(F(f)) \circ \eta_X$, which is true by the very definition of F we have given, these two results are in fact equal, meaning the diagram commutes. Therefore $\varphi : \text{Hom}_{\mathbf{C}}(F(-), -) \rightarrow \text{Hom}_{\mathbf{D}}(-, G(-))$ is natural in each of its arguments, and the functors F and G make up an adjoint pair. \square

This is a useful result indeed. To illustrate its utility, we will apply it to the example of abelianization from my last post. Recall that for each group G , the abelianization G^{ab} together with the natural surjection $\eta_G : G \rightarrow G/G^{\text{ab}}$ is a universal arrow from G to the inclusion functor $I : \mathbf{Ab} \rightarrow \mathbf{Grp}$. The theorem above proves that there is a functor $A : \mathbf{Grp} \rightarrow \mathbf{Ab}$ which is left adjoint to I . On objects, this functor is determined by $A(G) = G^{\text{ab}}$. For an arrow $f : G \rightarrow H$ in \mathbf{Grp} , $A(f)$ is the unique arrow such that the following diagram commutes:

$$\begin{array}{ccc}
G^{\text{ab}} & G \xrightarrow{\eta_G} & I(G^{\text{ab}}) \\
A(f) \downarrow & f \downarrow & \downarrow I(A(f)) \\
H^{\text{ab}} & H \xrightarrow{\eta_H} & I(H^{\text{ab}})
\end{array} \tag{9}$$

This functor is called the abelianization functor, and it is left adjoint to the inclusion functor. You can follow this same construction with any fully-defined set of universal arrows; the ones associated with the tensor product in the first post in this series are another good example. The key is that the universal construction must be possible for *all* objects in the codomain of the relevant functor; if there were a single group G which did not have an abelianization, we would not be able to build a left adjoint to I in this fashion.

Now, there are many more properties of adjoint functors we could discuss here. For example, the theorem we have just proven has a dual: given a functor $F : \mathbf{D} \rightarrow \mathbf{C}$ and a complete set of universal arrows *from* F , one can construct a functor $G : \mathbf{C} \rightarrow \mathbf{D}$ which is right adjoint to F . There is also a reverse sort of theorem; if we have an adjunction $F \dashv G$, we can construct natural transformations $\eta : 1_{\mathbf{D}} \rightarrow G \circ F$ and $\epsilon : F \circ G \rightarrow 1_{\mathbf{C}}$ with universal components, which are called the “unit” and “counit” of the adjunction. While these results are important and are quite interesting in their own right, proving them would be quite a bit more than what I want to do in this post. Instead, there is one very significant property of adjunctions which I do want to prove thoroughly, as it will be instrumental in the fourth (and final) post in this series.

Theorem 2. *Let $F : \mathbf{D} \rightarrow \mathbf{C}$, $G : \mathbf{C} \rightarrow \mathbf{D}$, and $\varphi : \text{Hom}_{\mathbf{C}}(F(-), -) \rightarrow \text{Hom}_{\mathbf{D}}(-, G(-))$ be an adjunction. Then G (the right adjoint) preserves limits, and dually F (the left adjoint) preserves colimits.*

Proof. We will show that G preserves limits; the proof that F preserves colimits is similar. To do this, let $D : \mathbf{J} \rightarrow \mathbf{C}$ be a diagram which has a limit $(L, \lambda_j : L \rightarrow D_j)$. We must show that the cone $(G(L), G(\lambda_j) : G(L) \rightarrow GD_j)$ is a limiting cone to $G \circ D$. In order to do this, let $(C, c_j : C \rightarrow GD_j)$ be an arbitrary cone to $G \circ D$. Using the isomorphism φ in the reverse direction, we arrive at $(F(C), \varphi^{-1}(c_j) : F(C) \rightarrow D_j)$; we claim first that this is a cone to D . To show this, let $\alpha : i \rightarrow j$ be an arbitrary arrow in \mathbf{J} - we must show that $D_\alpha \circ \varphi^{-1}(c_i) = \varphi^{-1}(c_j)$. To do so, we use the naturality of φ^{-1} , which gives the following commutative diagram:

$$\begin{array}{ccc}
\text{Hom}_{\mathbf{C}}(F(C), D_i) & \xleftarrow{\varphi^{-1}} & \text{Hom}_{\mathbf{D}}(C, GD_i) \\
\text{Hom}_{\mathbf{C}}(F(C), D_\alpha) \downarrow & & \downarrow \text{Hom}_{\mathbf{D}}(C, GD_\alpha) \\
\text{Hom}_{\mathbf{C}}(F(C), D_j) & \xleftarrow{\varphi^{-1}} & \text{Hom}_{\mathbf{D}}(C, GD_j)
\end{array} \tag{10}$$

Since $c_i \in \text{Hom}_{\mathbf{D}}(C, GD_i)$, we can trace it through each path of this diagram to get the equality $D_\alpha \circ \varphi^{-1}(c_i) = \varphi^{-1}(GD_\alpha \circ c_i)$. Then since (C, c_j) is a cone to $G \circ D$, $GD_\alpha \circ c_i = c_j$, giving us the equality $D_\alpha \circ \varphi^{-1}(c_i) = \varphi^{-1}(c_j)$ as desired. Now, because $(F(C), \varphi^{-1}(c_j))$ is a cone to D , and (L, λ_j) is a limit for D , there is a *unique* arrow $f : F(C) \rightarrow L$ such that for each $j \in \text{ob}(\mathbf{J})$, $\lambda_j \circ f = \varphi^{-1}(c_j)$. Now consider the arrow $\varphi(f) : C \rightarrow G(L)$. We claim that for each $j \in \text{ob}(\mathbf{J})$, $G(\lambda_j) \circ \varphi(f) = c_j$, recalling that our final goal is to produce a unique arrow with this property. To show this, first note that the naturality of φ means the following diagram commutes.

$$\begin{array}{ccc} \text{Hom}_{\mathbf{C}}(F(C), F(C)) & \xrightarrow{\varphi} & \text{Hom}_{\mathbf{C}}(C, G(F(C))) \\ \text{Hom}_{\mathbf{C}}(F(C), f) \downarrow & & \downarrow \text{Hom}_{\mathbf{D}}(C, G(f)) \\ \text{Hom}_{\mathbf{C}}(F(C), L) & \xrightarrow{\varphi} & \text{Hom}_{\mathbf{D}}(C, G(L)) \end{array} \quad (11)$$

Tracing the path of $1_{F(C)}$ through this diagram, we determine that $G(f) \circ \varphi(1_{F(C)}) = \varphi(f)$. This in turn means that $G(\lambda_j) \circ \varphi(f) = G(\lambda_j \circ f) \circ \varphi(1_{F(C)}) = G(\varphi^{-1}(c_j)) \circ \varphi(1_{F(C)})$. We can then use the naturality of φ again to show that the following diagram commutes for any $j \in \text{ob}(\mathbf{J})$:

$$\begin{array}{ccc} \text{Hom}_{\mathbf{C}}(F(C), F(C)) & \xrightarrow{\varphi} & \text{Hom}_{\mathbf{C}}(C, G(F(C))) \\ \text{Hom}_{\mathbf{C}}(F(C), \varphi^{-1}(c_j)) \downarrow & & \downarrow \text{Hom}_{\mathbf{D}}(C, G(\varphi^{-1}(c_j))) \\ \text{Hom}_{\mathbf{C}}(F(C), D_j) & \xrightarrow{\varphi} & \text{Hom}_{\mathbf{D}}(C, GD_j) \end{array} \quad (12)$$

Tracing the path of $1_{F(C)}$ through *this* diagram, we find that $G(\varphi^{-1}(c_j)) \circ \varphi(1_{F(C)}) = c_j$, completing the proof that $G(\lambda_j) \circ \varphi(f) = c_j$. The last step of our argument requires demonstrating that $\varphi(f)$ is the *only* arrow from C to $G(L)$ which commutes with the relevant cones. To demonstrate this, let $g' : C \rightarrow G(L)$ be an arrow such that for all $j \in \text{ob}(\mathbf{J})$, $G(\lambda_j) \circ g' = c_j$. Because φ is a bijection, there is an arrow $g : F(C) \rightarrow L$ such that $\varphi(g) = g'$. Additionally, the naturality of φ means that the following diagram must commute for any $j \in \text{ob}(\mathbf{J})$:

$$\begin{array}{ccc} \text{Hom}_{\mathbf{C}}(F(C), L) & \xleftarrow{\varphi^{-1}} & \text{Hom}_{\mathbf{D}}(C, G(L)) \\ \text{Hom}_{\mathbf{C}}(F(C), \lambda_j) \downarrow & & \downarrow \text{Hom}_{\mathbf{D}}(C, G(\lambda_j)) \\ \text{Hom}_{\mathbf{C}}(F(C), D_j) & \xleftarrow{\varphi^{-1}} & \text{Hom}_{\mathbf{D}}(C, GD_j) \end{array} \quad (13)$$

Tracing the paths of g' through this diagram, we determine that $\lambda_j \circ \varphi^{-1}(g') = \varphi^{-1}(G(\lambda_j) \circ g')$. By definition $G(\lambda_j) \circ g' = c_j$, and since $\varphi(g) = g'$, $\varphi^{-1}(g') = g$. Thus the previous equality may be rewritten as $\lambda_j \circ g = \varphi^{-1}(c_j)$. But since (L, λ_j) is a limit for D , there is exactly one arrow

with this property, namely f , so $g = f$. Then because φ is a bijection, $\varphi(f) = \varphi(g) = g'$, meaning $\varphi(f)$ has the desired uniqueness property. This completes the proof that G preserves limits. \square

This theorem is incredibly powerful because, as we have already seen, adjunctions seem to crop up everywhere in mathematics once you know where to look. For example, we already know that the diagonal functor preserves products (and all other limits) without even having to check it explicitly! And in light of theorem 1, all it takes to produce an adjunction is a complete set of universal arrows to a particular functor, after which one will be left with a pair of functors, one preserving limits and the other colimits.

That's all I have time for this time. Next time, in the fourth and final entry in this series, we'll focus back in on the abelianization and inclusion functors $A : \mathbf{Grp} \rightarrow \mathbf{Ab}$ and $I : \mathbf{Ab} \rightarrow \mathbf{Grp}$. Since we have already shown $A \dashv I$, we know A preserves colimits and I preserves limits. To drive home why all this abstract nonsense is powerful, I'll use this fact to prove a fascinating property in group theory that (in my estimation) would be very difficult to derive without categorical machinery. See you then!