Algebraic Geometry

Maxwell Jiang

Fall 2022

Abstract

Notes and solutions to exercises from Vakil's text.

Contents

1 Category Theory

1.1 Basic notions

Define categories, functors, natural transformations, etc.

Exercise 1.1.1. Show that functors preserve isomorphisms. Deduce that if $F_1, F_2 \colon \mathcal{A} \to \mathcal{B}$ and $G \colon \mathcal{B} \to \mathcal{C}$ are functors such that $F_1 \simeq F_2$ are naturally isomorphic, then $G \circ F_1 \simeq G \circ F_2$.

Solution. If $f \colon A \to A'$ is an isomorphism with inverse $f' \colon A' \to A$, then $F(f) \circ F(f') = F(f \circ f') = F(1_{A'}) = 1_{F(A')}$. Similarly $F(f') \circ F(f) = 1_{F(A)}$. Thus $F(f)$ is an isomorphism.

For the second statement, let $\eta \colon F_1 \simeq F_2$ denote the natural isomorphism. Apply G to the naturality square and note that the maps $\{G(\eta_A) \colon GF_1(A) \to GF_2(A)\}$ are isomorphisms by the previous part. □

We use $\mathrm{Hom}_{\mathcal{C}}(A, B)$ to denote the morphisms $A \to B$ in a category $\mathcal{C}$. Throughout, we work with **locally small** categories, i.e. $\mathrm{Hom}_{\mathcal{C}}(A, B)$ is always a set.

A **bifunctor** is a functor F defined on a product category $\mathcal{A} \times \mathcal{B}$. The most important example is $\mathrm{Hom} \colon \mathcal{C}^{\mathrm{op}} \times \mathcal{C} \to \mathbf{Set}$. Note that bifunctoriality is a stronger condition than "functoriality in each variable": given $A \to A'$ and $B \to B'$, bifunctoriality requires the following diagram to commute:

$$\begin{array}{ccc} F(A,B) & \longrightarrow & F(A',B) \\ \downarrow & & \downarrow \\ F(A,B') & \longrightarrow & F(A',B') \end{array}$$

A **subcategory** of $\mathcal{C}$ is formed by taking some of the objects and some of the morphisms of $\mathcal{C}$ such that identities are included and composition is respected. A functor $F \colon \mathcal{C}_1 \to \mathcal{C}_2$ **faithful** (resp. **full**) if all hom-set maps $\mathrm{Hom}_{\mathcal{C}_1}(A, B) \xrightarrow{F} \mathrm{Hom}_{\mathcal{C}_2}(FA, FB)$ are injective (resp. surjective). Note that we care more about morphisms than objects. We can view a subcategory as a faithful "inclusion functor" $\iota \colon \mathcal{A} \to \mathcal{B}$.

Exercise 1.1.2. Show that, in the category of finite-dimensional vector spaces, there is a natural isomorphism

$$V \cong (V^\vee)^\vee.$$

Solution. The isomorphism is given by the evaluation map

$$\mathrm{ev} \colon v \mapsto (f \mapsto f(v)).$$

This is an isomorphism when V is finite-dimensional. To check naturality, suppose $f \colon V \to W$ is a linear map. We need to show that the following diagram commutes:

$$\begin{array}{ccc} V & \xrightarrow{f} & W \\ \downarrow{\scriptstyle \mathrm{ev}} & & \downarrow{\scriptstyle \mathrm{ev}} \\ (V^\vee)^\vee & \xrightarrow{f^{**}} & (W^\vee)^\vee \end{array}$$

The induced map f^{**} is the map $\varphi \mapsto (\varphi \circ (g \mapsto g \circ f))$. If $v \in V$ is any vector, then following the diagram in either way yields the element $g \mapsto g(f(v))$. So the diagram commutes. □

A **monic** morphism (or **monomorphism**) is a map f such that $f \circ \alpha = f \circ \beta$ implies $\alpha = \beta$. Equivalently, the induced map on hom-sets is injective. Dually, an **epic** map (or **epimorphism**) is a map g such that $\alpha \circ g = \beta \circ g$ implies $\alpha = \beta$.

A **split monomorphism** is a map that has a left inverse; a **split epimorphism** is a map that has a right inverse. These are stronger conditions. They show up in the statement of the so-called *splitting lemma*.

1.2 Universal properties and constructions

Exercise 1.2.1. Show that any two initial objects are uniquely isomorphic. Similarly, any two final objects are uniquely isomorphic.

Solution. Suppose X and Y are initial. There's a unique $f\colon X \to Y$ and a unique $g\colon Y \to X$. The composition $fg\colon Y \to Y$ is unique, so it must be the identity 1_Y. Similarly $gf = 1_X$ and so $X \cong Y$ are uniquely isomorphic. For final objects the proof is the same. $\square$

In general, objects with universal properties are often initial or final in some auxiliary category.

1.2.1 Localization

We start with rings. Let A be a ring and $S \subset A$ a multiplicative subset. Define the **localization** $S^{-1}A$ as the set $(A \times S)/\sim$, where $(a_1, s_1) \sim (a_2, s_2)$ if and only if there exists $s \in S$ with $s(a_1 s - a_2 s_2) = 0$. The extra s term ensures that $\sim$ is transitive. We write a/s for the equivalence class $[(a, s)]$. Addition and multiplication are defined as they are for fraction fields. Note that $S^{-1}A = 0$ if $0 \in S$.

Exercise 1.2.2. Show that the canonical map $A \to S^{-1}A$ given by $a \mapsto a/1$ is injective if and only if S contains no zero-divisors of A.

Solution. We have $a/1 = 0/1$ if and only if $sa = 0$ for some $s \in S$. $\square$

Exercise 1.2.3. Check that $A \to S^{-1}A$ has the following universal property: it is initial with respect to A-algebras $A \to B$ for which S is mapped into $B^\times$.

Solution. We aim to find a unique map of A-algebras $\tilde{\varphi}\colon S^{-1}A \to B$, i.e.

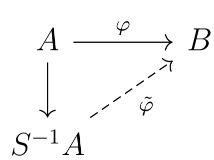

Any $\tilde{\varphi}$ satisying the commutative diagram needs to satisfy $\tilde{\varphi}(a/s)\varphi(s) = \tilde{\varphi}(a/s)\tilde{\varphi}(s) = \tilde{\varphi}(a) = \varphi(a)$, hence $\tilde{\varphi}(a/s) = \varphi(a)\varphi(s)^{-1}$ is forced. We check that this is well-defined: if $a_1/s_1 = a_2/s_2$, write $s(a_1 s_2 - a_2 s_1) = 0$ and apply φ. Cancel $\varphi(s)$ since it's a unit and rearrange to get $\varphi(a_1)\varphi(s_1)^{-1} = \varphi(a_2)\varphi(s_2)^{-1}$, as needed. $\square$

Consider the full subcategory $\mathbf{Mod}_{A,S}$ of $\mathbf{Mod}_A$ whose objects are A-modules M for which multiplication by $s \in S$ is invertible, i.e. the map $\mu\colon A \to \operatorname{End}_A(M)$ carries S into $\operatorname{Aut}_A(M)$.

Exercise 1.2.4. There is an equivalence of categories between $\mathbf{Mod}_{A,S}$ and $\mathbf{Mod}_{S^{-1}A}$.

Solution. Not going to be super precise, but here's the basic idea.

(i) Given an $S^{-1}A$-module M, we convert it to an A-module via restriction of scalars, i.e. $am := (a/1)m$. Conversely, suppose M is an A-module. Let $\overline{\operatorname{End}_{\mathbb{Z}}(M)}$ denote the (commutative) subring of $\operatorname{End}_{\mathbb{Z}}(M)$ generated by $\mu(A)$ and μ_s^{-1} for all $s \in S$. Now, by the universal property of Exercise 1.2.3, the map μ factors (uniquely) through a map $\tilde{\mu}$ as in the diagram below:

$$\begin{array}{ccc} A & \xrightarrow{\mu} & \overline{\operatorname{End}_{\mathbb{Z}}(M)} \\ \downarrow & \nearrow_{\tilde{\mu}} & \\ S^{-1}A & & \end{array}$$

In particular, we obtain an $S^{-1}A$-module structure on M compatible with the A-module structure. Explicitly, $(a/s)m = (\mu_a \circ \mu_s^{-1})(m)$.

(ii) Given any $S^{-1}A$-module morphism $f\colon M \to N$, it is clear that f also defines an A-module morphism $M \to N$, where the A-module structures are the ones prescribed by restriction of scalars as in (i). Conversely, suppose $f\colon M \to N$ is an A-module morphism. Then one can check that f respects the induced $S^{-1}A$-module structures on M and N using the explicit characterization given in (i).

□

Let's localize modules. Let M be an A-module. We'll define $S^{-1}M$ by universal property: $M \to S^{-1}M$ is initial among A-module maps $M \to N$ where $N \in \mathbf{Mod}_{A,S}$. To be precise, $M \to S^{-1}M$ is initial in the category whose objects are A-module maps $M \to N$ where $N \in \mathbf{Mod}_{A,S}$ and whose morphisms are given by A-module maps $N_1 \to N_2$ compatible with the respective maps from M. By Exericise 1.2.1, this is unique up to unique isomorphism.

The explicit construction of $S^{-1}M$ is just like localization of rings: take $M \times S$ modulo $(m_1, s_1) \sim (m_2, s_2)$ whenever $s(m_1 s_2 - m_2 s_1) = 0$ for some $s \in S$, etc. The canonical map $M \to S^{-1}M$ sends $m \mapsto m/1$. The A-module structure on $S^{-1}M$ is given by $a(m/s) := (am)/s$. By Exercise 1.2.4, this extends (uniquely) to an $S^{-1}A$-module structure. We'll see later that localizing modules is equivalent to extension of scalars via tensor products.

Exercise 1.2.5. Show that localization commutes with direct sums. That is, there's a natural isomorphism of $S^{-1}A$-modules

$$S^{-1}\left(\bigoplus M_\lambda\right) \cong \bigoplus S^{-1}M_\lambda.$$

Solution. Set $M = \bigoplus M_\lambda$. For each λ, consider the composition $M_\lambda \hookrightarrow M \mapsto S^{-1}M$. By the universal property, this factors uniquely through an $S^{-1}A$-module map $h_\lambda \colon S^{-1}M_\lambda \to S^{-1}M$. We show that $(S^{-1}M, \{h_\lambda\})$ satisfies the universal property of the direct sum $\bigoplus S^{-1}M_\lambda$. Let N be any $S^{-1}A$-module and suppose we have maps $\{f_\lambda \colon S^{-1}M_\lambda \to N\}_\lambda$. We wish to show that there is a unique $f \colon S^{-1}M \to N$ making the following diagram commutes for every λ.

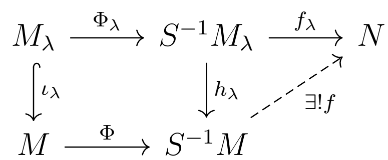

Consider the composition $g_\lambda = f_\lambda \circ \Phi_\lambda$. By the universal property of the direct sum M, the maps $\{g_\lambda\}$ assemble into a unique map $g \colon M \to N$ with $g_\lambda = g \circ \iota_\lambda$ for all λ. Now, the universal property of $S^{-1}M$ implies that $g = f \circ \Phi$ for a unique $f \colon S^{-1}M \to N$. We claim that this is the f we seek. Observe that

$$f_\lambda \circ \Phi_\lambda = g \circ \iota_\lambda = f \circ \Phi \circ \iota_\lambda = f \circ h_\lambda \circ \Phi_\lambda.$$

Now, the universal property of $S^{-1}M_\lambda$ implies that $f_\lambda = f \circ h_\lambda$, as needed. For uniqueness, notice that any valid f must satisfy $g = f \circ \Phi$, and we determined earlier that f is the unique map satisfying this property.

Naturality can be checked, but I don't feel like working out the details. □

Note. The above categorial solution is a little formal. Intuitively, the isomorphism depends on the "finiteness" of the direct sum (any vector has finitely many nonzero components). This means that for any element of $\bigoplus S^{-1}M_\lambda$, we can find a "common denominator", which yields an element of $S^{-1}M$. On the other hand, it is not true that localization commutes with arbitrary products.

1.2.2 Tensor products

Given A-modules M and N, their **tensor product** is an A-module $M \otimes_A N$ and an A-bilinear map $\pi \colon M \times N \to M \otimes_A N$ that is initial among such objects. Formally, for any A-module P, every bilinear map $b \colon M \times N \to P$ factors uniquely through π.

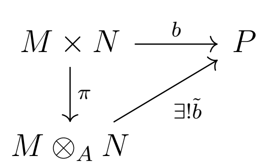

The tensor product is unique up to unique isomorphism. Explicitly, we take $M \otimes_A N$ to be the free A-module generated by $M \times N$ quotiented by the bilinearity relations. The map π is the composition of the canonical inclusion into the free module followed by the quotient.

Exercise 1.2.6. Show that the preceding construction for $M \otimes_A N$ satisfies the universal property of the tensor product.

Solution. We use the notation as above and let $A[M \times N]$ denote the free A-module on $M \times N$. By construction, $M \otimes_A N$ is generated by elements of the form $m \otimes n$, and any valid $\tilde{b}$ must satisfy $\tilde{b}(m \otimes n) = b(m, n)$. It follows that $\tilde{b}$ is unique if it exists.

We now construct $\tilde{b}$. By universal property of $A[M \times N]$, the map $b\colon M \times N \to P$ extends uniquely to an A-linear map $b\colon A[M \times N] \to P$. (This does not rely on bilinearity of b.) Let $q\colon A[M \times N] \to M \otimes_A N$ denote the quotient map. It suffices to show that $\ker q \subset \ker b$, and it will follow from the universal property of the quotient that b descends to a map $\tilde{b}$ with the desired properties.

$$\begin{array}{ccc} A[M \times N] & \xrightarrow{\ b\ } & P \\ \Big\downarrow q & \overset{\tilde{b}}{\nearrow} & \\ M \otimes_A N. & & \end{array}$$

At this point, we just check that every element of $A[M \times N]$ of the form $a(m, n) - (am, n)$, $a(m, n) - (m, an)$, $(m_1 + m_2, n) - (m_1, n) - (m_2, n)$, and $(m, n_1 + n_2) - (m, n_1) - (m, n_2)$ lies in $\ker b$ using bilinearity of the original map $b\colon M \times N \to P$. □

We denote $\pi(m, n)$ by $m \otimes n$ and call such elements *pure tensors*. One fact we deduce from the explicit construction of the tensor product is that $M \otimes_A N$ is generated by pure tensors. This is not apparent from the categorical definition, but it holds true for any tensor product of M and N since they are all isomorphic. It's clear from bilinearity that the A-module structure on $M \otimes_A N$ is given by $a(m \otimes n) = am \otimes n = m \otimes an$, etc.

Typically, the universal property of $M \otimes_A N$ is used to check that maps on $M \otimes_A N$ are well-defined. To specify a map $f\colon M \otimes_A N \to P$, we simply declare its values on $m \otimes n$ and then check that $(m, n) \mapsto m \otimes n \mapsto f(m \otimes n)$ is bilinear. Then f exists and is unique.

Exercise 1.2.7. Show that $(-) \otimes_A (-)$ defines a bifunctor $\mathbf{Mod}_A \times \mathbf{Mod}_A \to \mathbf{Mod}_A$.

Solution. Consider maps $f\colon M \to M'$ and $g\colon N \to N'$. The composition $\pi' \circ (f \times g)$ from $M \times N \to M' \times N' \to M' \otimes_A N'$ is bilinear, and we obtain a unique map $(f \times g)_*$ making the following diagram commute:

$$\begin{array}{ccc} M \times N & & \\ \Big\downarrow \pi & \overset{\pi' \circ (f \times g)}{\searrow} & \\ M \otimes_A N & \overset{\exists !}{\dashrightarrow} & M' \otimes_A N' \end{array}$$

Clearly $(-)_*$ preserves identities. It also respects composition. Consider the commutative diagram

$$\begin{array}{ccccc} M \times N & \xrightarrow{f \times g} & M' \times N' & \xrightarrow{f' \times g'} & M'' \times N'' \\ \Big\downarrow \pi & & \Big\downarrow \pi' & & \Big\downarrow \pi'' \\ M \otimes_A N & \dashrightarrow & M' \otimes_A N' & \dashrightarrow & M'' \otimes_A N'' \end{array}$$

The map $(f'f \times g'g)_*$ is the unique map h such that $\pi'' \circ (f'f \times g'g) = h \circ \pi$. Meanwhile, we have

$$\begin{aligned}(f' \times g')_* \circ (f \times g)_* \circ \pi &= (f' \times g')_* \circ \pi' \circ (f \times g) \\ &= \pi'' \circ (f' \times g') \circ (f \times g) \\ &= \pi'' \circ (f'f \times g'g).\end{aligned}$$

Thus $(f'f \times g'g)_* = (f' \times g')_* \circ (f \times g)_*$, as needed. □

Exercise 1.2.8 (Extension of scalars)**.** Let M be an A-module and $A \to B$ an A-algebra. Then $B \otimes_A M$ has a B-module structure that is compatible with its A-module structure via restriction of scalars. Moreover, $B \otimes_A (-)$ is a functor $\mathbf{Mod}_A \to \mathbf{Mod}_B$.

Solution. For any $b \in B$, consider the map $B \xrightarrow{b} B$ given by multiplication by b. Clearly this is a map of A-modules, so functoriality of the tensor product in the first slot yields an induced map $B \otimes_A M \xrightarrow{b_*} B \otimes_A M$ which defines scalar multiplication by b on $B \otimes_A M$. (We'll sometimes use $(\cdot)$ for scalar multiplication to emphasize that it is by an element of B.) Explicitly, we have $b \cdot (b' \otimes m) = bb' \otimes m$. Associativity of this scalar multiplication follows from the fact that functors respect composition. Compatibility with the A-module structure on $B \otimes_A M$ follows (tautologically, almost) since the A-module structure on B is *defined* by restriction of scalars. Indeed, writing $\varphi \colon A \to B$ for the structure map, we have

$$a(b \otimes m) = (\varphi(a)b) \otimes m = \varphi(a) \cdot (b \otimes m).$$

Functoriality of the tensor product in the second slot implies that $B \otimes_A (-)$ is a functor from A-modules to A-modules. In particular, each $f \colon M \to M'$ induces an A-module map $f_* \colon B \otimes_A M \to B \otimes_A M'$. It suffices to check that f_* respects the B-module structures. This amounts to commutativity of the following diagram for any $b \in B$.

$$\begin{array}{ccc} B \otimes_A M & \xrightarrow{f_*} & B \otimes_A M' \\ \downarrow b_* & & \downarrow b_* \\ B \otimes_A M & \xrightarrow{f_*} & B \otimes_A M' \end{array}$$

This follows from bifunctoriality of the tensor product. □

Exercise 1.2.9. Suppose $A \to B$ and $A \to C$ are A-algebras. Then $B \otimes_A C$ has a natural ring structure.

Solution. Let $b \in B$ and $c \in C$ be arbitrary. The maps $B \xrightarrow{b} B$ and $C \xrightarrow{c} C$ induce an A-module map $B \otimes_A C \xrightarrow{(b,c)_*} B \otimes_A C$. Explicitly, this map is given by $b' \otimes c' \mapsto bb' \otimes cc'$. We can then define multiplication by $(b \otimes c)(b' \otimes c') = bb' \otimes cc'$, and it is clear that this makes $B \otimes_A C$ into a (commutative, unital) ring.

I don't actually know what natural means in this context. □

Exercise 1.2.10 (Localization is extension of scalars)**.** Let S be a multiplicative subset of A and M an A-module. Then there is a natural isomorphism of $S^{-1}A$-modules

$$(S^{-1}A) \otimes_A M \cong S^{-1}M.$$

Solution. Consider the map $M \to (S^{-1}A) \otimes_A M$ given by $m \mapsto 1 \otimes m$. We show that this satisfies the universal property of $S^{-1}M$. Let N be any $S^{-1}A$-module and consider an A-module map $f\colon M \to N$. We wish to show that there is a unique $S^{-1}A$-module map $\tilde{f}$ making the diagram commute.

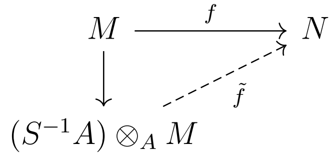

Any valid $\tilde{f}$ must satisfy $\tilde{f}(1 \otimes m) = f(m)$ and thus is determined on all pure tensors, hence unique if it exists. Define $\tilde{f}$ by $c \otimes m \mapsto cf(m)$ for all $c \in S^{-1}A$ and $m \in M$. Checking the appropriate A-bilinearity conditions shows that $\tilde{f}$ is well-defined as an A-module map. By Exercise 1.2.4, it is also a $S^{-1}A$-module map, as needed.

Here's a sketch for naturality. Consider a map $f\colon M \to N$. Since $S^{-1}M$ and $(S^{-1}A)M \otimes_A M$ both satisfy universal properties, the induced maps $S^{-1}M \to S^{-1}N$ and $(S^{-1}A) \otimes_A M \to (S^{-1}A) \otimes_A N$ are unique for their respective commuting squares. We can then argue that a certain pair of maps must coincide, which yields the desired naturality. □

Exercises 1.2.10 and 1.2.5 also give natural isomorphisms if both sides are considered as A-modules; just apply the restriction of scalars functor and use Exercise 1.1.1.

Exercise 1.2.11. Show that tensor products commute with arbitrary direct sums. That is, there's a natural isomorphism of A-modules

$$M \otimes_A \left(\bigoplus N_\lambda\right) \cong \bigoplus M \otimes_A N_\lambda.$$

Solution. Define $\Phi\colon M \otimes_A (\bigoplus N_\lambda) \to \bigoplus M \otimes_A N_\lambda$ by $m \otimes (n_\lambda) \mapsto (m \otimes n_\lambda)$.

To define the inverse map, first consider the maps $\psi_\mu\colon M \otimes_A N_\mu \to M \otimes_A (\bigoplus N_\lambda)$ given by $m \otimes n \mapsto m \otimes (\delta_{\lambda\mu} n)$, where $(\delta_{\lambda\mu} n)$ is the vector with n in the μ component and zeroes elsewhere. The universal property of the direct sum yields a map $\Psi\colon \bigoplus M \otimes_A N_\lambda \to M \otimes_A (\bigoplus N_\lambda)$ given by $\Psi = \sum_\lambda \psi_\lambda$.

It's easy to see that Φ and Ψ are inverse homomorphisms, so the desired isomorphism follows. As for naturality, I don't feel like checking the details. □

1.2.3 Fiber products and pullback squares

Given maps $X \xrightarrow{\alpha} Z$ and $Y \xrightarrow{\beta} Z$ in any category, a **fibered product** is an object $X \times_Z Y$ equipped with maps π_X and π_Y to X and Y such that the following **pullback square** commutes:

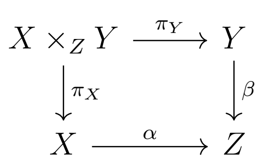

Moreover, $(X \times_Z Y, \pi_X, \pi_Y)$ is final among such triples. Formally, if W is an object with maps π'_X and π'_Y to X and Y such that $\alpha \circ \pi'_X = \beta \circ \pi'_Y$, then π'_X and π'_Y both factor through a unique map $W \to X \times_Z Y$.

In **Set**, the fibered product is the subset of the product $X \times Y$ consisting of pairs (x, y) with $\alpha(x) = \beta(y)$.

Exercise 1.2.12. Suppose Z is a final object. Assuming they exist, show that $X \times_Z Y$ and $X \times Y$ are uniquely isomorphic.

Solution. Draw the diagram:

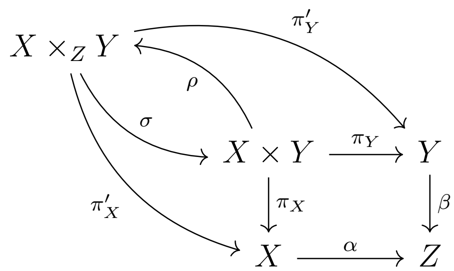

The universal property of the product yields a unique map σ such that $\pi'_X = \pi_X \circ \sigma$ and $\pi'_Y = \pi_Y \circ \sigma$. Meanwhile, finality of Z implies that the maps $\alpha \circ \pi_X$ and $\beta \circ \pi_Y$ must equal the unique map $X \times Y \to Z$, hence the universal property of the fibered product yields a unique map ρ such that $\pi_X = \pi'_X \circ \rho$ and $\pi_Y = \pi'_Y \circ \rho$. The standard argument now shows that ρ and σ are inverses, as needed. Moreover, the preceding discussion implies that ρ and σ are the *unique* isomorphism pair between $(X \times Y, \pi_X, \pi_Y)$ and $(X \times_Z Y, \pi'_X, \pi'_Y)$. $\square$

Exercise 1.2.13. Given morphisms $X_1 \to Y, X_2 \to Y$, and $Y \to Z$, show that there is a natural morphism $X_1 \times_Y X \to X_1 \times_Z X_2$, assuming that both fibered products exist.

Solution. The maps $X_1 \times_Y X_2 \xrightarrow{\pi_1} X_1 \to Y \to Z$ and $X_1 \times_Y X_2 \xrightarrow{\pi_2} X_2 \to Y \to Z$ agree, so the universal property of $X_1 \times_Z X_2$ yields the desired map. $\square$

Exercise 1.2.14 (Magic diagram)**.** With notation as in the previous exercise, show that

$$\begin{array}{ccc} X_1 \times_Y X_2 & \longrightarrow & X_1 \times_Z X_2 \\ \downarrow & & \downarrow \\ Y & \longrightarrow & Y \times_Z Y \end{array}$$

is a pullback square. (Assume all relevant fibered products exist.)

Solution. First, let's clarify where the maps in the magic diagram come from and why they commute. Just take a close look at the following commutative diagram:

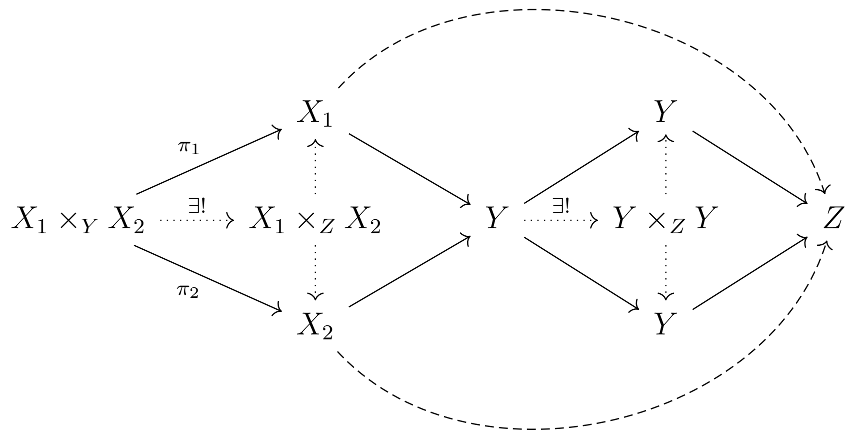

Suppose maps $T \to Y$ and $T \to X_1 \times_Z X_2$ are given such that $f\colon T \to Y \to Y \times_Z Y$ and $g\colon T \to X_1 \times_Z X_2 \to Y \times_Z Y$ agree. Check that the following diagrams commute:

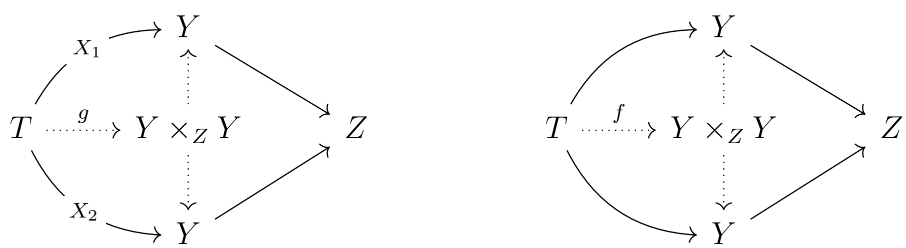

The maps $T \to X_i \to Y \to Z$ on the left are the ones obtained by composition with $T \to X_1 \times_Z X_2$, and the maps $T \to Y$ on the right are both the given map $T \to Y$. Note that $f = g$ holds if and only if $T \to X_1 \to Y$ and $T \to X_2 \to Y$ both equal $T \to Y$, which is equivalent to the existence of a unique map $h\colon T \to X_1 \times_Y X_2$ "making everything commute". □

We define the **coproduct** and **fibered coproduct** by reversing all arrows in the definitions of product and fibered product, respecitvely. For example, the coproduct in **Set** is disjoint union. We use **pushout square** to denote the defining commutative square of the fibered coproduct.

Exercise 1.2.15. Recall from Exercise 1.2.9 that if $\sigma\colon A \to B$ and $\tau\colon A \to C$ are A-algebras, the tensor product $B \otimes_A C$ inherits a natural ring structure. Show that, equipped with the maps $B \to B \otimes_A C$ and $C \to B \otimes_A C$ given by $b \mapsto b \otimes 1$ and $c \mapsto 1 \otimes c$, this construction is the fibered coproduct of $A \to B$ and $A \to C$ in **CRing**.

Solution. Owing to the A-module structure on $B \otimes_A C$, we have $\sigma(a) \otimes 1 = a(1 \otimes 1) = 1 \otimes \tau(a)$

for any $a \in A$. Thus the square in the diagram below commutes.

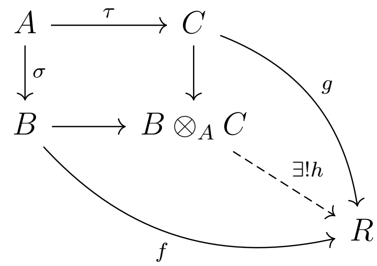

To check that it is a pushout square, let R be a ring and $f\colon B \to R$, $g\colon C \to R$ ring maps with $f \circ \sigma = g \circ \tau$. (Note that this endows R with the structure of an A-algebra.) Any ring map $h\colon B \otimes_A C \to R$ making the diagram commute is uniquely determined: it must satisfy $h(b \otimes 1) = f(b)$ and $h(1 \otimes c) = g(c)$, meaning it must satisfy $h(b \otimes c) = f(b)g(c)$, and is thus determined on all of $B \otimes_A C$. To show that h exists, we simply note that $(b, c) \mapsto f(b)g(c)$ is A-bilinear and thus $h\colon b \otimes c \mapsto f(b)g(c)$ is a well-defined A-module map. It's clear that h defined as such is a ring map (in fact, an A-algebra map), and we're done. $\square$

Note. We can equivalently interpret $B \otimes_A C$ as the coproduct of B and C in $\mathbf{Alg}_A$.

1.2.4 Yoneda lemma

Exercise 1.2.16 (Yoneda's lemma)**.** Suppose F is a covariant functor $\mathcal{C} \to \mathbf{Set}$ and $A \in \mathcal{C}$ is an object. Then there is a bijection between $F(A)$ and the set of natural transformations $\eta\colon \operatorname{Hom}_{\mathcal{C}}(A, -) \to F$.

Solution. Let η be such a natural transformation. Naturality implies that, for every morphism $f\colon A \to B$, the following diagram commutes.

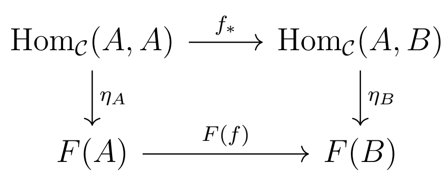

Define $\theta(\eta) := \eta_A(1_A)$. Note that the square implies that $\eta_B(f) = F(f)(\theta(\eta))$. As B and f were arbitrary, it follows that η is completely determined by $\theta(\eta)$. In particular, θ is injective as a map from natural transformations to $F(A)$. Moreover, given an element $x \in F(A)$, it is easy to see that setting $\eta_B(f) := F(f)(x)$ for each $B \in \mathcal{C}$ and $f \in \operatorname{Hom}_{\mathcal{C}}(A, B)$ defines a natural transformation $\eta\colon \operatorname{Hom}_{\mathcal{C}}(A, -) \to F$. In particular, θ is surjective. Thus θ yields the desired bijection. $\square$

There is a dual formulation of Yoneda's lemma where F is instead a *contravariant* functor and the natural transformations η take $\operatorname{Hom}_{\mathcal{C}}(-, A)$ to F. The proof is nearly identitcal. We use the dual to describe a special case of the lemma. Consider the **functor category** of $\mathcal{C}$, denoted $\mathcal{F}(\mathcal{C}^{\mathrm{op}})$, whose objects are contravariant functors $\mathcal{C} \to \mathbf{Set}$ and whose morphisms are natural transformations between said functors. There is a (covariant) functor $h_\bullet\colon \mathcal{C} \to \mathcal{F}(\mathcal{C}^{\mathrm{op}})$ sending A to $\operatorname{Hom}_{\mathcal{C}}(-, A)$, called the **Yoneda embedding**.

For a given $B \in \mathcal{C}$, if we let F denote the (contravariant) functor $\mathrm{Hom}_{\mathcal{C}}(-, B)$, then the Yoneda lemma yields a bijection between $F(A) = \mathrm{Hom}_{\mathcal{C}}(A, B)$ and the set of natural transformations from $\mathrm{Hom}_{\mathcal{C}}(-, A)$ to $F = \mathrm{Hom}_{\mathcal{C}}(-, B)$. In other words, the Yoneda embedding induces a bijection between hom-sets, and is thus a fully faithful functor. (Hence "embedding".)

Of course, there is a dual formulation of the Yoneda embedding, which states that the (contravariant) functor $h^\bullet$ embeds $\mathcal{C}$ into the category of covariant functors $\mathcal{C} \to \mathbf{Set}$.

I guess the slogan here is: morphisms $A \to B$ are the same as natural transformations $\mathrm{Hom}_{\mathcal{C}}(-, A) \to \mathrm{Hom}_{\mathcal{C}}(-, B)$, and the same as natural transformations $\mathrm{Hom}_{\mathcal{C}}(B, -) \to \mathrm{Hom}_{\mathcal{C}}(A, -)$.

1.3 Limits and colimits

A **small** category is a category whose objects form a set. (Think posets.) Let $\mathcal{I}$ be a small category and $A\colon \mathcal{I} \to \mathcal{C}$ a functor. We call A a **diagram indexed by $\mathcal{I}$**. Intuitively, the data of A is a commutative diagram in $\mathcal{C}$ whose "shape" is given by $\mathcal{I}$.

The **limit** of the diagram, denoted $\varprojlim A_i$, is an object L equipped with morphisms $f_i\colon L \to A_i$ for each $i \in \mathcal{I}$ such that, for any morphism $m\colon j \to k$ in $\mathcal{I}$, the following diagram commutes:

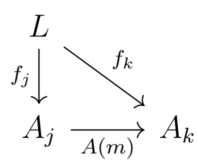

We require the limit to be final with respect to this property, and it follows that $\varprojlim A_i$ is unique up to unique isomorphism.

For example, if $\mathcal{I}$ has only the identity morphisms, then the limit $\varprojlim A_i$ is the product $\prod A_i$. If $\mathcal{I}$ has three objects $\{1, 2, 3\}$ whose only nonidentity morphisms are $1 \to 3$ and $2 \to 3$, the $\varprojlim A_i$ is the fibered product $A_1 \times_{A_3} A_2$.

Exercise 1.3.1. Suppose $\mathcal{I}$ is a poset with an initial object e. Show that the limit of any diagram indexed by $\mathcal{I}$ exists.

Solution. We show that the limit is A_e, where for each $i \in \mathcal{I}$ the map $A_e \to A_i$ is the one induced by the unique map $e \to i$. By functoriality, all the required commuting triangles are satisfied.

Suppose L is another object with maps $g_i\colon L \to A_i$ for every $i \in \mathcal{I}$ satisfying the commuting triangles. Note that these maps include a map $g_e\colon L \to A_e$. We wish to show that g_e is the unique map making the following diagram commute for all $m\colon j \to k$ in $\mathcal{I}$:

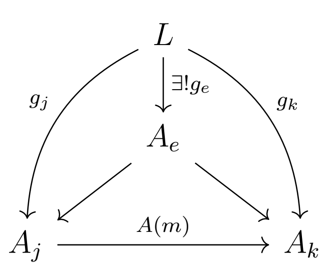

By assumption g_e works. For uniqueness, simply take m to be the identity $A_e \to A_e$. Then trivially any valid $L \to A_e$ equals g_e, as desired. $\square$

Exercise 1.3.2. Show that

$$\varprojlim A_i = \{(a_i)_{i \in \mathcal{I}} \mid A(m)\colon a_j \mapsto a_k \text{ for all } m\colon j \to k\}$$

equipped with the coordinate projections $\pi_j\colon \varprojlim A_i \to A_j$ gives an explicit construction of limits in **Set**.

Solution. Given any $m\colon j \to k$ and any element $a = (a_i)_{i \in \mathcal{I}}$ in the limit, we have

$$A(m)(\pi_j(a)) = A(m)(a_j) = a_k = \pi_k(a),$$

so the proposed limit satisfies all commuting triangles. Suppose $(L, \{f_i\colon L \to A_i\})$ also satisfies all commuting triangles. The function $L \to \varprojlim A_i$ given by $\ell \mapsto (f_i(\ell))_{i \in \mathcal{I}}$ makes everything commute; it is also easy to see that is unique. $\square$

The construction in Exercise 1.3.2 works equally well for categories consisting of "sets and functions with additional structure", like $\mathbf{Mod}_A$ and **CRing**. Think of the construction as the subset of the product $\prod A_i$ consisting of "chains" generated by the maps in the diagram. Example: consider the ring $\mathbb{Z}_p$ of p-adic integers, defined as the limit $\varprojlim \mathbb{Z}/p^i$.

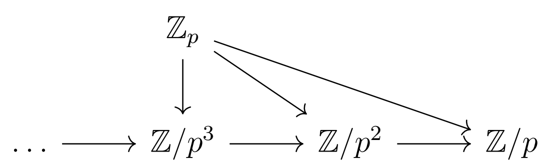

Here, the index category $\mathcal{I}$ is the opposite cateogry of the poset $(\mathbb{N}, \le)$, and the maps $Z/p^{i+1} \to Z/p^i$ are the obvious reduction-modulo-p^i maps. In light of Exercise 1.3.2, one often defines $\mathbb{Z}_p$ as the ring whose elements are formal series $a_0 + a_1 p + a_2 p^2 + \dots$ where $0 \le a_i < p$. Each truncation $a_0 + a_1 p + \cdots + a_i p^i$ determines an element of $\mathbb{Z}/p^{i+1}$.

Let's now talk about **colimits**. It's notated $\varinjlim A_i$. You can guess the definition. The coproduct is the colimit when $\mathcal{I}$ has no nontrivial morphisms.

Exercise 1.3.3. Interpret $\mathbb{Q} = \varinjlim \frac{1}{n}\mathbb{Z}$.

Solution. For now, we'll do it in **Set**. The diagram in question is indexed by the poset $\mathbb{N}$ ordered by divisibility, and the maps $\frac{1}{n}\mathbb{Z} \to \frac{1}{m}\mathbb{Z}$ for $n \mid m$ as well as $\frac{1}{n}\mathbb{Z} \to \mathbb{Q}$ are all inclusions.

To show that $\mathbb{Q}$ is initial, suppose T is a set equipped with maps $f_n\colon \frac{1}{n}\mathbb{Z} \to T$ satisfying all commuting triangles. We seek a unique $f\colon \mathbb{Q} \to T$ making everything commute. Given any $q \in \mathbb{Q}$, write $q = a/b$ in lowest terms and set $f(q) = f_b(a/b)$. Note that this is forced, so f is unique. To see that f satisfies all commuting triangles, suppose $a/b = a'/b'$. Then $b \mid b'$, and we have the following diagram, which commutes except possibly for the triangle on the right.

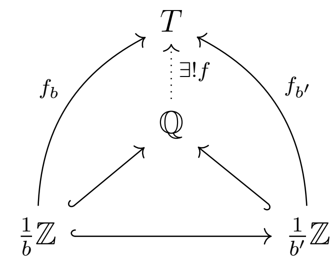

We have $f_{b'}(a'/b') = f_b(a/b) = f(\iota_{b'}(a'/b'))$, so the triangle on the right commutes, too. $\square$

We introduce a nice class of index categories for which there is a simple description of the colimit. A nonempty category $\mathcal{I}$ is **filtered** if

(i) For any $i, j \in \mathcal{I}$, there is an object $k \in \mathcal{I}$ and morphisms $i \to k$ and $j \to k$,

(ii) For any parallel morphisms $m_1, m_2 \colon i \to j$, there is a morphism $\pi \colon j \to k$ for which $\pi \circ m_1 = \pi \circ m_2$.

Note that if $i \to j$ and $i \to k$ are morphisms in a filtered category, then there exist morphisms $j \to \ell$ and $k \to \ell$ that "complete the square": $i \to j \to \ell$ equals $i \to k \to \ell$.

Exercise 1.3.4. Suppose $\mathcal{I}$ is filtered. Show that any diagram $\{A_i\}$ in **Set** indexed by $\mathcal{I}$ has the following colimit:

$$\varinjlim A_i = \left\{ (a_i, i) \in \coprod_{i \in \mathcal{I}} A_i \right\} / \sim, \tag{1.3.1}$$

where $\sim$ identifies (a_i, i) and (a_j, j) if and only if the diagram contains morphisms $f \colon A_i \to A_k$ and $g \colon A_j \to A_k$ such that $f(a_i) = g(a_j)$. The maps $A_j \to \varinjlim A_i$ are the obvious ones.

Solution. We first check that $\sim$ is actually an equivalence relation: reflexivity follows from identity morphisms and symmetry is evident. For transitivity, suppose $(a_i, i) \sim (a_j, j)$ and $(a_j, j) \sim (a_k, k)$. Then there exist $f \colon A_i \to A_m$ and $g \colon A_j \to A_m$ for which $f(a_i) = g(a_j) = a_m$, and $f' \colon A_j \to A_\ell$ and $g' \colon A_k \to A_\ell$ for which $f'(a_j) = g'(a_k) = a_\ell$. Pick morphisms $f'' \colon A_m \to A_n$ and $g'' \colon A_\ell \to A_n$ such that $f'' \circ g = g'' \circ f'$. It follows that $(f'' \circ f')(a_i) = (g'' \circ g')(a_k)$, so $(a_i, i) \sim (a_k, k)$, as needed.

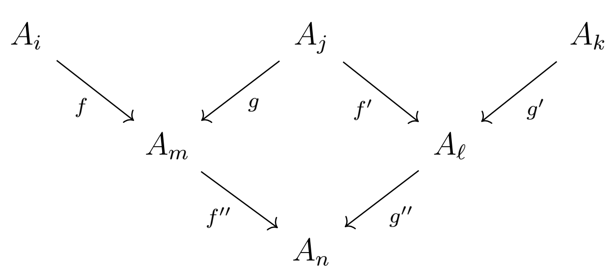

We now show that $\varinjlim A_i$ is initial. Suppose T is a set equipped with maps $f_i \colon A_i \to T$ satisfying all commuting triangles. We wish to find a unique $f \colon \varinjlim A_i \to T$ making everything commute. Note that every element $x \in \varinjlim A_i$ is the image of some a_j under the map $A_j \to \varinjlim A_i$. Thus we are forced to set $f(x) = f_j(a_j)$. It remains to show that this is well-defined; in other words, that $f_j(a_j) = f_k(a_k)$ whenever $(a_j, j) \sim (a_k, k)$. Pick maps $g\colon A_j \to A_m$ and $h\colon A_k \to A_m$, belonging to the diagram, such that $g(a_j) = h(a_k)$. By assumption, the following diagram commutes:

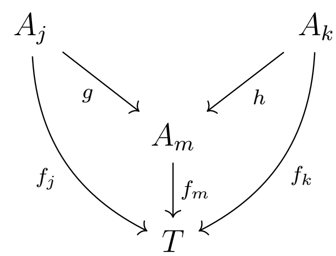

So $f_j(a_j) = f_m(g(a_j)) = f_m(h(a_k)) = f_k(a_k)$, as needed. □

The way to think about the equivalence relation $\sim$ in (1.3.1) is that it's the equivalence relation "generated by the morphisms of the diagram", identifying an element of A_i with its images.

The construction in Exercise 1.3.4 also works in $\mathbf{Mod}_A$. More precisely, if $\mathcal{I}$ is filtered, the underlying set of $\varinjlim M_i$ is given by (1.3.1), and addition is defined as follows: for any $m_i \in M_i$ and $m_j \in M_j$, find morphisms $M_i \to M_k$ and $M_j \to M_k$ belonging to the diagram and sum the images of m_i, m_j to obtain an element $m_k \in M_k$. We then declare

$$[(m_i, i)] + [(m_j, j)] := [(m_k, k)].$$

For scalar multiplication, we set $a[(m_i, i)] = [(am_i, i)]$.

Exercise 1.3.5. Show that the above declarations turn $\varprojlim M_i$ as defined in (1.3.1) into the colimit of the diagram in $\mathbf{Mod}_A$.

Solution. We show that addition is well-defined. First we show that it is independent of the choice of M_k. In what follows, all morphisms are assumed to belong to the diagram. Suppose $f\colon M_i \to M_k$, $g\colon M_j \to M_k$, $f'\colon M_i \to M_{k'}$, and $g'\colon M_j \to M_{k'}$. Pick $\alpha\colon M_k \to M_\ell$ and $\beta\colon M_{k'} \to M_\ell$ so that $\alpha \circ f = \beta \circ f'$, and then pick $\gamma\colon M_\ell \to M_n$ such that $\gamma \circ \alpha \circ g = \gamma \circ \beta \circ g'$. Then $\gamma(\alpha(f(m_i) + g(m_j))) = \gamma(\beta(f'(m_i) + g'(m_j)))$ and so $[(f(m_i) + g(m_j), k)] = [(f'(m_i) + g'(m_j), k')]$. To show that addition is independent of the representatives m_i and m_j, complete some more squares. There are more things to check but they're straightforward. □

Exercise 1.3.6 (Localization as a filtered colimit)**.** Let A be a domain and $S \subset A$ multiplicative. Then, in the category of A-modules,

$$S^{-1}A = \varinjlim \frac{1}{s}A.$$

Here, the limit is taken over $s \in S$, and we view $\frac{1}{s}A$ as a submodule of $\operatorname{Frac}(A)$. (Note that $\frac{1}{s}A$ might not be a ring!) We interpret $\frac{1}{0}A$ as all of $\operatorname{Frac}(A)$. The morphisms of the diagram are the inclusions $\frac{1}{s}A \to \frac{1}{s'}A$ for $s \mid s'$, and the maps $\frac{1}{s}A \to S^{-1}A$ are the obvious ones.

Solution. Observe that the index category (namely, S with morphisms given by divisibility) is filtered: for any s_1, s_2, there are morphisms $s_1 \to s_1 s_2$ and $s_2 \to s_1 s_2$, and any parallel morphisms agree since there's at most one morphism between any two objects.

Checking that $S^{-1}A$ is the colimit follows much the same procedure as in Exercise 1.3.3. In fact, this exercise allows us to interpret Exercise 1.3.3 in $\mathbf{Mod}_{\mathbb{Z}}$ instead of **Set**. □

In set-like categories, we should liken limits to "intersections" and colimits to "unions" (and in certain categories, like posets that are power sets under inclusion, this can be made precise). Intuitively, an element of a limit is an element "belonging to all the objects", and an element of the colimit is a "representative" for an element in some object.

Let's now think about limits and colimits as functors. Suppose $\mathcal{C}$ is a category in which arbitrary limits and colimits exist. Let $\mathcal{F}(\mathcal{I}, \mathcal{C})$ denote the category of diagrams of shape $\mathcal{I}$ (i.e. functors from $\mathcal{I} \to \mathcal{C}$). The morphisms are natural transformations, i.e. collections of maps $(f_i)_{i\in\mathcal{I}}$ making everything commute. Then there are covariant functors

$$\varprojlim_{\mathcal{I}}, \varinjlim_{\mathcal{I}} \colon \mathcal{F}(\mathcal{I}, \mathcal{C}) \to \mathcal{C}$$

sending each diagram D to its limit and colimit, respectively. Here's a picture showing where the induced map of colimits comes from:

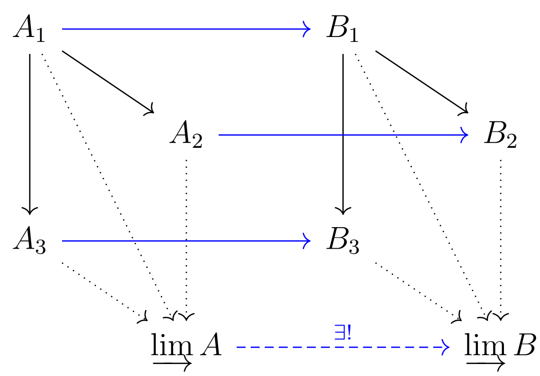

Exercise 1.3.7. Make sense of the statment "limits commute with limits" and prove it. Similarly, colimits commute with colimits.

Solution. Suppose $\mathcal{I}, \mathcal{J}$ are index categories, and we have a functor

$$D \colon \mathcal{I} \to \mathcal{F}(\mathcal{J}, \mathcal{C}).$$

Think of D as a "shape-$\mathcal{I}$ diagram of shape-$\mathcal{J}$ diagrams". The claim is that

$$\varprojlim_{\mathcal{I}} \left(\varprojlim_{\mathcal{J}} \circ D \right) = \varprojlim_{\mathcal{J}} \left(\varprojlim_{\mathcal{I}} D \right). \tag{1.3.2}$$

On the left side, the inner term is the composition of functors, and it sends each $i \in \mathcal{I}$ to the limit of its corresponding $\mathcal{J}$-shaped diagram; Doing this for all i yields a "$\mathcal{I}$-shaped diagram of limits", and we can then take the limit of these limits over $\mathcal{I}$.

On the right, we first "take a limit of diagrams", producing a diagram (and a collection of natural transformations) in the shape of $\mathcal{J}$, and we then take the limit of that diagram over $\mathcal{J}$.

We sketch a proof. First show that, for a given $j \in \mathcal{J}$, the object $\left(\varprojlim_{\mathcal{I}} D\right)_j$ is the limit of the diagram whose objects are $\{(D_i)_j \mid i \in \mathcal{I}\}$ and whose morphisms are the "jth components" of the natural transfomrations between the diagrams D_i.

Let L denote the object on the right side of Equation (1.3.2). Then L comes equipped with maps into the "limit of diagrams", which compose with the maps comprising the natural maps from that diagram into each of the diagrams D_i. These maps induce maps from L into $\varprojlim_{\mathcal{J}} D_i$ for each i, and everything commutes. It remains to show that L equipped with the latter maps is final. Suppose $\tilde{L}$ is another. Then by composing stuff we get maps from $\tilde{L}$ into each of the D_i, and using finality of each $\left(\varprojlim_{\mathcal{I}} D\right)_j$ we get maps from $\tilde{L}$ into the diagram $\varprojlim_{\mathcal{I}} D$. These then induce a map into L, and we're done.

For colimits, just do the same thing with some arrows reversed. (Note that the colimit functor is still covariant, however.) □

Here's a picture of an example, with kernels (defined in a later section).

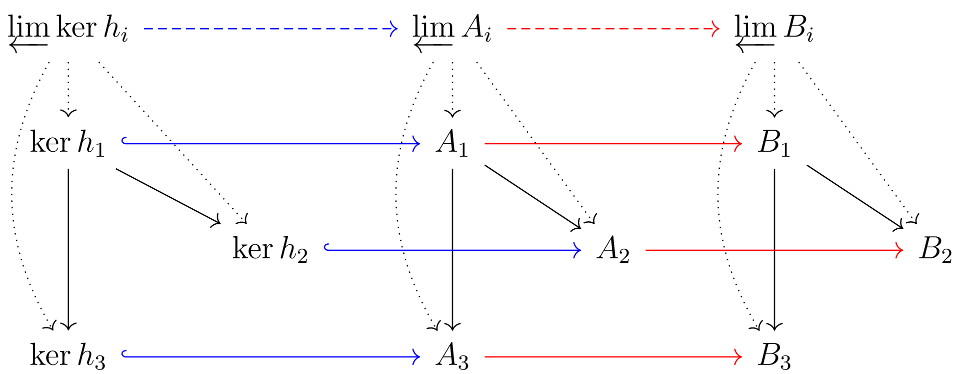

This comes from the interpretation of kernels as limits, though I've omitted a bunch of zero objects for clarity. The content of Exercise 1.3.7 is that the kernel on the top left object is actually the kernel of the map on the top right.

1.4 Adjoints

Given (covariant) functors $F\colon \mathcal{A} \to \mathcal{B}$ and $G\colon \mathcal{B} \to \mathcal{A}$, we say (F, G) is an **adjoint pair** if for all $A \in \mathcal{A}$ and $B \in \mathcal{B}$, there is a natural bijection

$$\tau_{AB}\colon \operatorname{Hom}_{\mathcal{B}}(F(A), B) \simeq \operatorname{Hom}_{\mathcal{A}}(A, G(B)).$$

Visually, the bijection is between the red and blue arrows in the diagram below:

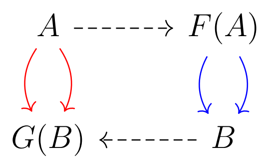

Note that that $\text{Hom}_{\mathcal{B}}(F(-),-)$ and $\text{Hom}_{\mathcal{A}}(-,G(-))$ are bifunctors $\mathcal{A}^{\text{op}} \times \mathcal{B} \to \mathbf{Set}$. By naturality, we mean that τ is a natural isomorphism between them. Explicitly, for any morphisms $A' \to A$ and $B \to B'$, we have the following commuting square in **Set**.

$$\begin{array}{ccc} \text{Hom}_{\mathcal{B}}(F(A),B) & \xrightarrow{\tau_{AB}} & \text{Hom}_{\mathcal{A}}(A,G(B)) \\ \downarrow & & \downarrow \\ \text{Hom}_{\mathcal{B}}(F(A'),B') & \xrightarrow{\tau_{A'B'}} & \text{Hom}_{\mathcal{A}}(A',G(B')) \end{array}$$

An adjunction says that "the data of a morphism $F(A) \to B$ is the same as the data of a morphism $A \to G(B)$". The example to keep in mind is the adjuction between the free and forgetful functors (say, between **Set** and **Grp**). Indeed, given a set S, any group homomorphism $F[S] \to H$ defined on the free group generated by S is uniquely specified by a function $S \to H$ of sets.

Exercise 1.4.1 (Units and counits)**.** Suppose (F, G) is an adjoint pair. For each A, there is a natural morphism $\eta_A : A \to GF(A)$, called the **unit** of the adjunction, with the following property. For any morphism $g : F(A) \to B$, the corresponding $\tau_{AB}(g) : A \to G(B)$ is given by the composition

$$A \xrightarrow{\eta_A} GF(A) \xrightarrow{G(g)} G(B).$$

Formulate the dual statement.

Solution. We set $\eta_A := \tau_{A,F(A)}(1_{F(A)})$. We wish to show that the left triangle in the following diagram commutes:

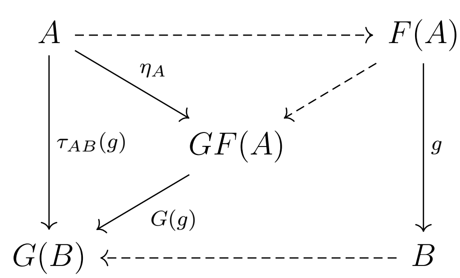

This is simply a consequence of the naturality of τ:

$$\begin{array}{ccc} \text{Hom}_{\mathcal{B}}(F(A),B) & \xrightarrow{\tau_{AB}} & \text{Hom}_{\mathcal{A}}(A,G(B)) \\ {\scriptstyle G(g)_*}\uparrow & & {\scriptstyle G(g)_*}\uparrow \\ \text{Hom}_{\mathcal{B}}(F(A),F(A)) & \xrightarrow{\tau_{A,F(A)}} & \text{Hom}_{\mathcal{A}}(A,GF(A)) \end{array}$$

I don't feel like checking naturality or formulating the dual. $\square$

The unit in the adjunction between the free and forgetful functors is the "canonical inclusion" of the generating set into the underlying set of the free group it generates. Dually, the counit is the "introducing relations" group homomorphism $F[G] \to G$ that sends $g \mapsto g$.

Exercise 1.4.2 (Currying isomorphism)**.** Let M, N, P be A-modules. Describe a natural isomorphism

$$\operatorname{Hom}_A(M \otimes_A N, P) \cong \operatorname{Hom}_A(M, \operatorname{Hom}_A(N, P)).$$

Deduce that $(-) \otimes_A N$ and $\operatorname{Hom}_A(N, -)$ are adjoint functors.

Solution. In the left-to-right direction, the isomorphism τ is given by

$$\sigma \mapsto (m \mapsto (n \mapsto \sigma(m \otimes n))). \tag{1.4.1}$$

Using the universal property of the tensor product, we can construct an inverse map

$$\theta \mapsto (m \otimes n \mapsto \theta(m)(n)).$$

We check naturality of τ. Suppose $f\colon M' \to M$ and $g\colon P \to P'$ are A-module maps. We observe the effect of the induced maps on both sides of (1.4.1).

(i) The left side is sent to $g \circ \sigma \circ (f \otimes 1)$.

(ii) The right side is sent to $(m \mapsto (n \mapsto g(\sigma(f(m) \otimes n))))$.

Evidently, τ sends (i) to (ii), so the induced maps respect τ. □

If we endow the hom-sets in $\mathbf{Mod}_A$ with A-module structures, we see that the currying isomorphism is in fact a natural isomorphism in $\mathbf{Mod}_A$ and not just **Set**.

Exercise 1.4.3 (Restriction and extension of scalars are adjoints)**.** Let $A \to B$ be an A-algebra. Let $M \to M_A$ denote the restriction-of-scalars functor $\mathbf{Mod}_B \to \mathbf{Mod}_A$, and recall from Exercise 1.2.8 the extension-of-scalars functor $\mathbf{Mod}_A \to \mathbf{Mod}_B$ given by $B \otimes_A (-)$.

Show that $M \to M_A$ is right-adjoint to $B \otimes_A (-)$.

Solution. Let M be an A-module and N a B-module. We seek a bijection τ from the B module maps $B \otimes_A M \to N$ to A-module maps $M \to N_A$.

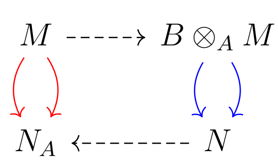

Suppose $g\colon B \otimes_A M \to N$ is a B-module map. Define $\tau(g)$ to be the map $m \mapsto g(1 \otimes m)$. For concreteness, let's check that $\tau(g)$ actually defines an A-module map. Writing $\varphi\colon A \to B$ for the structure map, we have

$$\begin{aligned}
\tau(g)(am) &= g(1 \otimes am) \\
&= g(a(1 \otimes m)) \\
&= g(\varphi(a) \cdot (1 \otimes m)) \\
&= \varphi(a) \cdot g(1 \otimes m) \\
&= ag(1 \otimes m),
\end{aligned}$$

as needed. For the inverse, suppose $f\colon M \to N_A$ is an A-module map. Define $\rho(f)$ to be the map $b \otimes m \mapsto b \cdot f(m)$. Check that $\rho(f)$ is a B-module map.

We check that τ and ρ are inverses. The map $\rho(\tau(g))$ sends $b \otimes m$ to $b \cdot g(1 \otimes m) = g(b \otimes m)$. The map $\tau(\rho(f))$ sends m to $(b \otimes m' \mapsto b \cdot f(m'))(1 \otimes m) = f(m)$.

Too lazy to check naturality. □

Note. In both the free/forgetful adjunction and in Exercise 1.4.3, the functor that "forgets structure" is right-adjoint to the functor that "adds structure".

Exercise 1.4.4. Show that right adjoints commute with limits and left adjoints commute with colimits.

Solution. Let $\mathcal{C}_1, \mathcal{C}_2$ be categories and suppose (F, G) is an adjoint pair between them. Let $A = \{A_i\}_{i \in \mathcal{I}}$ be a diagram in $\mathcal{C}_2$. We wish to show that

$$G(\varprojlim A) = \varprojlim (GA).$$

Let $T \in \mathcal{C}_1$ be an object equipped with maps (red) into GA making everything commute. It suffices to show that there is a unique map $T \to G(\varprojlim A)$ making everything commute.

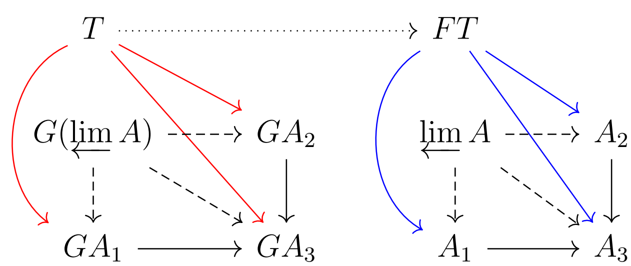

Now we use the adjoint property to port everything over to $\mathcal{C}_2$. The red maps are in bijection with the blue maps, and everything commutes by naturality. Then it's enough to show that there's a unique map $FT \to \varprojlim A$ making everything commute in $\mathcal{C}_2$, but this simply follows from the definition of the limit.

The dual property for left adjoints and colimits is proved similarly. □

For example, Exercise 1.4.4 implies that colimits commute with tensor products.

1.5 Abelian categories

A **zero object** is an object that's initial and final. Clearly, any two zero objects are uniquely isomorphic.

Exercise 1.5.1. In a category $\mathcal{C}$ with a zero object, there is a unique *zero map* $0\colon A \to 0 \to B$ between any two objects A and B. Composition with the zero map yields the zero map.

Solution. For uniqueness, suppose we have two zero maps, as shown. By definition, all maps in the diagram exist and are unique, so the diagram commutes.

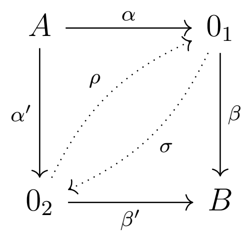

Then $\beta\alpha = \beta\rho\alpha' = \beta'\alpha'$, as desired. The second property is obvious. □

Suppose $\mathcal{C}$ has a zero object. The **kernel** of a map $f : A \to B$ is a an object $\ker(f)$ along with a map $\iota : \ker(f) \to A$ such that $f \circ \iota = 0$, and $(\ker(f), \iota)$ is final with respect to this property. Dually, a **cokernel** of f is an object $\operatorname{coker}(f)$ along with a map $\pi : C \to \operatorname{coker}(f)$ such that $\pi \circ f = 0$, and $(\operatorname{coker}(f), \pi)$ is initial with respect to this property. If f is a monomorphism, we may refer to $\operatorname{coker}(f)$ as a **quotient** and write B/A. Of course, kernels and cokernels are unique up to unique isomorphism. We'll abuse notation and say "kernel" or "cokernel" to refer to the object, the map, or both.

Exercise 1.5.2. Interpret kernels and cokernels as limits and colimits, respectively. Hint: the diagrams will have three objects, one of them zero...

Exercise 1.5.3. A kernel is monic. Dually, a cokernel is epic.

Solution. Suppose $g, h : C \to \ker(f)$ are parallel morphisms such that $\iota \circ g = \iota \circ h$, call this map α. Then $f \circ \alpha = 0$, hence α factors *uniquely* through a map to the kernel. It follows that $g = h$, as needed.

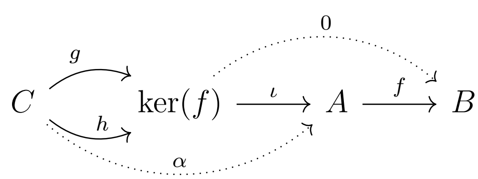

The dual is similar. □

Exercise 1.5.4. Consider maps $f : A \to B$ and $f : A \to C$. Show that the assertion $\ker(f) = \ker(g)$ makes sense. In other words, if there exists a map $k : K \to A$ that is a kernel of both f and g, then any kernel of f is a kernel of g and vice versa.

Solution. The existence of a common kernel implies that any map α into A satisfies $f \circ \alpha = 0$ if and only if $g \circ \alpha = 0$, as both are equivalent to α factoring through k. Now, say a little more stuff and finish. □

An **abelian category** is a category $\mathcal{C}$ with the following properties:

(i) Every hom-set $\operatorname{Hom}_{\mathcal{C}}(A, B)$ is equipped with the structure of an abelian group such that composition distributes over addition.

(ii) There exists a zero object in $\mathcal{C}$.

(iii) Finite products exist.

(iv) Kernels and cokernels exist.

(v) Every monomorphism is the kernel of its cokernel.

(vi) Every epimorphism is the cokernel of its kernel.

An **additive category** is one satisfying (i), (ii), (iii). An **additive functor** between additive categories is a functor that respects addition of maps (in other words, determines abelian group homomorphisms between hom-sets.)

Exercise 1.5.5. In an additive category, the additive identity $0_{AB} \in \mathrm{Hom}_{\mathcal{C}}(A, B)$ is the zero map $0\colon A \to B$.

Solution. Let $\beta \circ \alpha$ denote the zero map $A \xrightarrow{\alpha} 0 \xrightarrow{\beta} B$. By finality of 0, we have $\alpha + \alpha = \alpha$. Hence $\beta \circ \alpha = \beta \circ (\alpha + \alpha) = \beta \circ \alpha + \beta \circ \alpha$ and thus $\beta \circ \alpha = 0_{AB}$. □

Note that in an additive category, the endomorphisms of an object form a (possibly noncommutative, unital) ring. Recall that a ring is the zero ring if and only if $1 = 0$.

Exercise 1.5.6. In an additive category, an object X is a zero object if and only if $1_X = 0_X$ and deduce that additive functors send zero objects to zero objects.

Solution. If X is a zero object, then $\mathrm{End}_{\mathcal{C}}(X)$ consists of a single morphism, so $1_X = 0_X$. Conversely, suppose $1_X = 0_X$. For any morphism $f\colon X \to Y$, we have $f = f \circ 1_X = f \circ 0_X = 0$. Similarly, for any morphism $f\colon Y \to X$, we have $f = 1_X \circ f = 0_X \circ f = 0$. Thus X is a zero object.

An additive functor F preserves identities and zero maps. So if X is a zero object, then $1_{F(X)} = F(1_X) = F(0_X) = 0_{F(X)}$ so $F(X)$ is also a zero object. □

Exercise 1.5.7. In an additive category, a map f is monic if and only if $f \circ x = 0$ implies $x = 0$, if and only $\ker(f)$ is a zero object. Dual statement for epic maps and cokernels.

Solution. Easy. Note that additivity implies that any f with $\ker(f) = 0$ is monic; the other direction is true without the additive assumption. □

It's common practice to assume that functors between additive categories are additive, but we'll be explicit and specify each time.

Exercise 1.5.8. In an abelian category, monic and epic implies isomorphism.

Solution. By Exercise 1.5.7, a monic and epic map $f\colon A \to B$ has kernel and cokernel 0. By conditions (v) and (vi) in the definition of abelian category, f is a cokernel of $0 \to A$ and a kernel of $B \to 0$. Thus there exist unique morphisms taking the place of the dotted arrows below making the diagram commute.

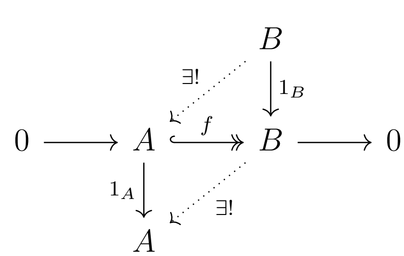

It follows that f has both a left inverse and a right inverse, hence f has an inverse and is thus an isomorphism. $\square$

Define the **image** of a map $f\colon A \to B$ by $\mathrm{im}(f) := \ker(\mathrm{coker}(f))$. By definition, images always exist in abelian categories. Check that images are unique up to unique isomorphism (easy, but not immediate from the corresponding fact for kernels and cokernels).

Exercise 1.5.9. In an abelian category, every map $f\colon A \to B$ factors uniquely through a map $\tilde{\pi}\colon A \to \mathrm{im}(f)$, and π is epic. Moreover, $\tilde{\pi}\colon A \to \mathrm{im}(f)$ is a cokernel of $\ker(f)$.

Solution. As a kernel, $\mathrm{im}(f)$ comes with a map $\tilde{\iota}\colon \mathrm{im}(f) \hookrightarrow B$. As $\pi \circ \tilde{f} = 0$, the map f factors uniquely through a map $\tilde{\pi}\colon A \to \mathrm{im}(f)$. We claim that the kernel $\ker(f) \hookrightarrow A$ of f is also a kernel of $\tilde{\pi}$. Indeed, we have $f \circ \iota = \tilde{\iota} \circ \tilde{\pi} \circ \iota = 0$ and monicniess of $\tilde{\iota}$ implies that $\tilde{\pi} \circ \iota = 0$. Moreover, if $g\colon C \to A$ is a map with $\tilde{\pi} \circ g = 0$ then $f \circ g = \tilde{\iota} \circ \tilde{\pi} \circ g = 0$, hence g factors uniquely through a map $h\colon C \to \ker(f)$.

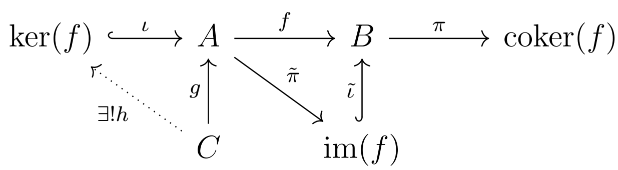

It remains to show that $\tilde{\pi}$ is epic; it will follow from the preceding claim along with condition (vi) in the definition of abelian category that $\tilde{\pi}$ is the cokernel of $\ker(f) \hookrightarrow A$. So, suppose $x\colon \mathrm{im}(f) \to X$ is a map with $x \circ \tilde{\pi} = 0$. We wish to show that $x = 0$. Let $k\colon \ker(x) \hookrightarrow \mathrm{im}(f)$ be its kernel, so that it suffices to show that k is epic.

As $x \circ \tilde{\pi} = 0$, there's a unique $\rho\colon A \to \ker(x)$ with $\tilde{\pi} = k \circ \rho$. Since $\tilde{\iota} \circ k$ is a composition of monic maps, it is monic, hence it's the kernel of a map $t\colon B \to T$ by condition (v). But now we have $t \circ f = t \circ \tilde{\iota} \circ k \circ \rho = 0$, hence t factors uniquely as $t = \sigma \circ \pi$ for some $\sigma\colon \mathrm{coker}(f) \to T$.

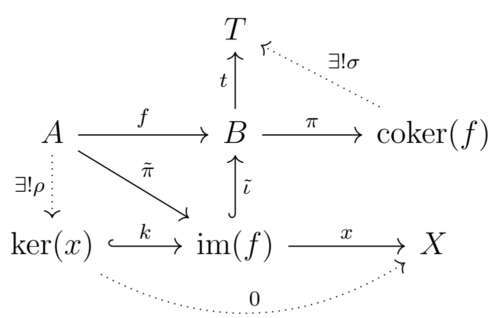

But now $t \circ \tilde{\iota} = \sigma \circ \pi \circ \tilde{\iota} = 0$. Recall that $\ker(x)$ is the kernel of t, so as a result $\tilde{\iota}$ factors through a map $\tilde{k}\colon \mathrm{im}(f) \to \ker(x)$. Then $\tilde{\iota} = \tilde{\iota} \circ k \circ \tilde{k}$, and by monicness of $\tilde{\iota}$ we have $k \circ \tilde{k} = 1$. In particular, k has a right inverse, hence is epic, and we're done. $\square$

Exercise 1.5.10. Consider the following diagram in an abelian category. Suppose the central square commutes. Show that there are unique morphisms in the place of the dotted arrows making the diagram commute.

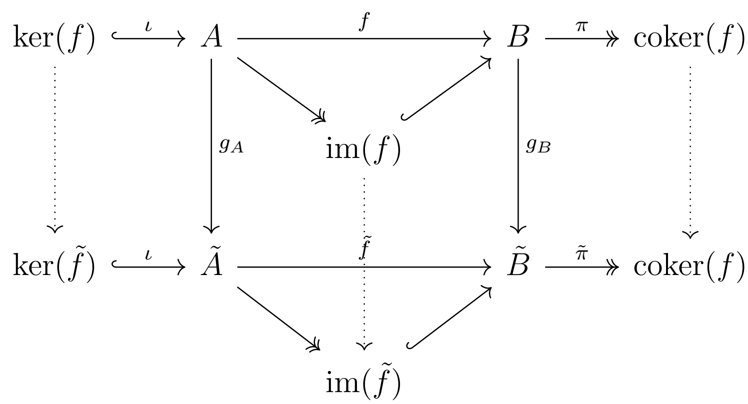

Solution. We start with kernels: we have $\tilde{f} \circ g_A \circ \iota = g_B \circ f \circ \iota = 0$, hence $g_A \circ \iota$ factors uniquely through a map to $\ker \tilde{f}$, as needed. Analogous thing for cokernels. By Exercise 1.5.9, the images $\mathrm{im}(f)$ and $\mathrm{im}(\tilde{f})$ are cokernels of $\ker(f)$ and $\ker(\tilde{f})$; hence the result for cokernels implies that there's a unique map $\mathrm{im}(f) \to \mathrm{im}(\tilde{f})$ making the left parallelogram, with A and $\mathrm{im}(\tilde{f})$ as corners, commute. One can show that any such map necessarily makes the right parallelogram, with $\mathrm{im}(\tilde{f})$ and B as corners, commute (hint: use monicness of the inclusion $\mathrm{im}(\tilde{B}) \hookrightarrow B$). So everything commutes and we're happy. □

Note. Exercise 1.5.10 is pretty useful. For example, consider two compositions $A \hookrightarrow B \to B/A$ and $\tilde{A} \hookrightarrow \tilde{B} \to \tilde{B}/\tilde{A}$. Compatible maps $A \to \tilde{A}$ and $B \to \tilde{B}$ induce a unique compatible map between the quotients.

Let $\mathcal{C}$ be an abelian category throughout. A **chain complex** in $\mathcal{C}$, denoted $C_\bullet$, is a diagram of the form $(\ldots \xrightarrow{\partial} C_n \xrightarrow{\partial} C_{n-1} \xrightarrow{\partial} \ldots)$ with $\partial^2 = 0$. Latter condition implies there's a canonical monomorphism $\mathrm{im}\, \partial_{n+1} \hookrightarrow \ker \partial_n$. We define **$n$th homology** $H_n(C_\bullet) := \ker \partial_n / \mathrm{im}\, \partial_{n+1}$. (For chain complexes with increasing indices, we write $H^n(C^\bullet)$ and call it *cohomology*). A complex is **exact** at C_n if $\mathrm{im}\, \partial_{n+1} \hookrightarrow \ker \partial_n$ is an isomorphism, equivalently $H_n(C_\bullet) = 0$, equivalently there is an object that is simultaneously an image of ∂_{n+1} and a kernel of ∂_n. Check that homology is a functor. (Use Exercise 1.5.10, twice!)

Let $\mathbf{Com}_\mathcal{C}$ denote the category of chain complexes over $\mathcal{C}$. The morphisms are chain maps.

Exercise 1.5.11. Show that $\mathbf{Com}_\mathcal{C}$ is an abelian category.

Solution. Not going to do all the details, just use the fact that $\mathcal{C}$ is abelian and tack things together to get the corresponding properties for complexes. It's easy to see that $\mathrm{Hom}_{\mathbf{Com}_\mathcal{C}}(C_\bullet, D_\bullet)$ is an abelian group (add chain maps componentwise). There's a zero object $(\cdots \to 0 \to 0 \to \ldots)$. Kernels always exist, namely, tack together the componentwise kernels using the induced maps of Exercise 1.5.10. (Can check that it satisfies $\partial^2 = 0$ using monicness of kernels.) Similar thing for cokernels. □

Exercise 1.5.12. Show that a covariant additive functor between abelian categories is left exact if and only if it preserves kernels. Dually, it is right exact if and only if it preserves cokernels.

For contravariant functors, the respective conditions should be that cokernels are sent to kernels and vice versa.

Solution. We'll just do the covariant case. This boils down to showing that $f\colon A \to B$ has kernel $\iota\colon K \to A$ if and only if $0 \to K \xrightarrow{\iota} A \xrightarrow{f} B$ is exact. (Observe that the functor sends 0 to 0 by Exercise 1.5.6, which is why additivity is needed.) By Exercise 1.5.7, the latter is equivalent to exactness at A and monicness of ι. If this holds, then $\ker(f) = \operatorname{im}(\iota) = \ker(\operatorname{coker}(\iota))$. But a monic map in an abelian category is the kernel of its cokernel, so ι works for the right side. By Exercise 1.5.4, this means ι also works as a kernel of f, as needed. Converse is similar. □

Define **left exact** and **right exact** functors between abelian categories (both covariant and contravariant – remember to flip the arrows for the latter). Define **exact** functors as those that are left exact and right exact.

Exercise 1.5.13. Show that (additive) exact functors preserve exactness. (That is, $A \to B \to C$ exact implies $FA \to FB \to FC$ exact.)

Solution. Show that exactness of $A \to B \to C$ is equivalent to $\ker(g)$ being a kernel of $\operatorname{coker}(f)$, then use Exercise 1.5.12. □

Exercise 1.5.14 (Exactness properties of $\otimes$ and Hom)**.** Show that

(i) Given an A-module N, the functor $(-) \otimes_A N$ is a right exact functor.

(ii) Localization of modules is exact.

(iii) Given an A-module N, the functors $\operatorname{Hom}_A(-, N)$ and $\operatorname{Hom}_A(N, -)$ are left exact functors from $\mathbf{Mod}_A$ to itself. In a general abelian category $\mathcal{C}$ and an object $X \in \mathcal{C}$, the functors $\operatorname{Hom}_{\mathcal{C}}(-, X)$ and $\operatorname{Hom}_{\mathcal{C}}(X, -)$ are left exact functors from $\mathcal{C}$ to $\mathbf{Ab}$.

Solution. We first note that all functors in question are additive, so Exercise 1.5.12 applies.

For (i): Suppose $M' \xrightarrow{f} M \xrightarrow{g} M'' \to 0$ is exact. The map $g \otimes 1$ induced by tensoring sends $m \otimes n$ to $g(m) \otimes n$, hence hits all generators of the codomain, hence is surjective. It remains to show exactness of

$$M' \otimes_A N \xrightarrow{f\otimes 1} M \otimes_A N \xrightarrow{g\otimes 1} M'' \otimes_A N.$$

It's clear that the above composition is 0, hence $g \otimes 1$ factors through a map $\tilde{g}$ from the quotient $(M \otimes_A N)/\operatorname{im}(f \otimes 1)$. It's enough to show that $\tilde{g}$ is an isomorphism. To do this, we can construct an explicit inverse. By surjectivity of g, we may choose a preimage $m \in g^{-1}(m'')$ for each $m'' \in M''$. Define the inverse map by $m''\otimes n \mapsto m \otimes n \pmod{\operatorname{im}(f\otimes 1)}$. To see that this is well-defined, suppose $m_1, m_2 \in M$ are such that $g(m_1) = g(m_2)$. Then $m_1 \otimes n - m_2 \otimes n = (m_1 - m_2) \otimes n \equiv 0 \pmod{\operatorname{im}(f \otimes 1)}$, where the last equality uses exactness of the original sequence.

For (ii): Let A be a ring, $S \subset A$ a multiplicative subset. Recall from Exercise 1.2.4 that localization of A-modules is equivalent to tensoring by $S^{-1}A$, and by (i) it suffices to show that localization is left exact. By Exercise 1.5.12, we just need to show that it preserves kernels. Consider $\ker(f) \hookrightarrow M \xrightarrow{f} M'$. Let's take the explicit constructions. We wish to show that the kernel of $f_*\colon S^{-1}M \to S^{-1}M'$ is precisely $\{m/s \mid m \in \ker f\}$. Note that every element of $S^{-1}M$ takes the form m/s. (If working with the tensor product interpretation, then every element of $(S^{-1}A) \otimes M$ is a finite combination of elements of the form $(a/s) \otimes m$, and the idea is that we can "put the fraction over a common denominator" to get it all into one term.) So, if $f(m)/s = 0$, then multiplying both sides by s yields $f(m) = 0$, as needed.

For (iii), we'll just show that $\mathrm{Hom}_A(-, N)$ is left exact, and the rest are similar. It suffices to show that it sends cokernels to kernels. Suppose $M' \xrightarrow{f} M \xrightarrow{g} M''$ where g is a cokernel of f. Any map $\sigma\colon M \to S$ with $\sigma \circ f = 0$ factors uniquely through g. Now consider the dualized sequence

$$\mathrm{Hom}_A(M', N) \xleftarrow{f^*} \mathrm{Hom}_A(M, N) \xleftarrow{g^*} \mathrm{Hom}_A(M'', N).$$

Then g^* is injective, and $f^* \circ g^* = 0$. The kernel of f^*, as a submodule of $\mathrm{Hom}_A(M, N)$, consists of those φ such that $\varphi \circ f = 0$. But we observed earlier that such φ must factor through g, equivalently they lie in the image of g^*. The result follows. $\square$

Note. It's possible to prove (i) using (iii) via the currying isomorphism. Also, here's an example of the tensor product failing to be left exact. Consider

$$0 \to \mathbb{Z} \xrightarrow{2} \mathbb{Z}.$$

Now tensor everything by $\mathbb{Z}/2\mathbb{Z}$. Kernels are not preserved: the second map becomes the zero map.

Exercise 1.5.15 (FHHF)**.** Suppose $C_\bullet$ is a chain complex over an abelian category $\mathcal{C}$. Let F be a covariant additive functor from $\mathcal{C}$ to another abelian category.

(i) If F is right exact, then there is a natural map $FH_\bullet \to H_\bullet F$.

(ii) If F is left exact, then there is a natural map in the opposite direction.

(iii) If F is exact, then maps in (i) and (ii) are inverses and thus F "commutes with homology".

Solution. Apply F to the complex $A_\bullet$ and consider the commuting diagram:

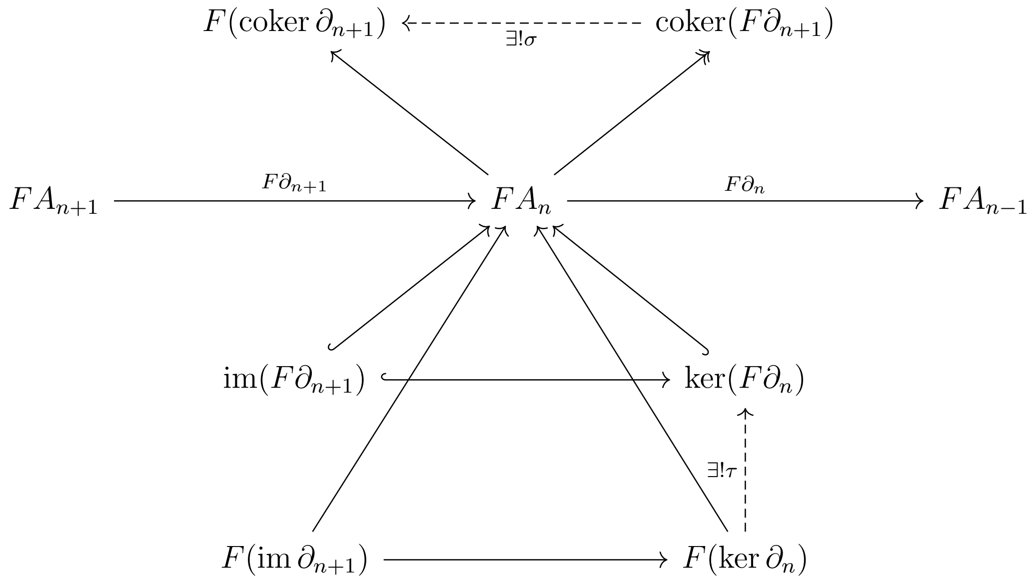

The maps σ and τ between the cokernels and kernels follow from universal properties. (For example, the composition $FA_{n+1} \to FA_n \to F(\operatorname{coker}\partial_{n+1})$ is zero before and thus after applying F.)

Suppose F is right exact. Since F preserves cokernels, $F(H_n(A_\bullet))$ is the cokernel of the bottom map in the above diagram. Additionally, the map σ is an isomorphism. It follows that the composition

$$F(\operatorname{im}\partial_{n+1}) \to FA_n \to \operatorname{coker}(F\partial_{n+1})$$

is zero, so it factors uniquely through a map $F(\operatorname{im}\partial_{n+1}) \to \operatorname{im}(F\partial_{n+1})$. We claim that this makes the bottom square commute; to see this, just compose both paths with the monomorphism $\ker(F\partial_n) \hookrightarrow FA_n$ and use the rest of the diagram to show that the results are equal. Now, apply Exercise 1.5.10 to get the map $F(H_n(A_\bullet)) \to H_n(FA_\bullet)$ between the cokernels, which is what we needed.

Next suppose F is left exact, i.e. preserves kernels. So this time τ is an isomorphism, and $F(\operatorname{im}\partial_{n+1}) \hookrightarrow FA_n$ is the kernel of $FA_n \to F(\operatorname{coker}\partial_{n+1})$. Since the composition $\operatorname{im}(F\partial_{n+1}) \hookrightarrow FA_n \to F(\operatorname{coker}\partial_{n+1})$ is zero (factor through σ), it factors uniquely through a map $\operatorname{im}(F\partial_{n+1}) \to F(\operatorname{im}\partial_{n+1})$. As before, we can show that the bottom square commutes (this time with downward vertical arrows), and we get a map between the cokernels. However, it is not necessarily the case that $F(H_n(A_\bullet))$ is a cokernel of the bottom row, so we get the actual desired map from a diagram like this:

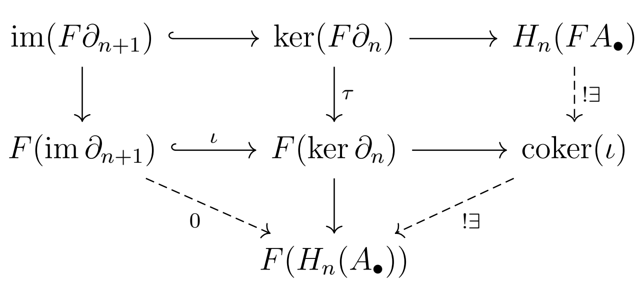

I'm lazy to show naturality and do (iii) but they seem intuitive enough. □

For a concrete example of Exercise 1.5.15, consider $\mathbf{Mod}_R$ and let F be the tensor product functor $(-) \otimes_R N$. Part (i) says that there is a natural map

$$H_n(A_\bullet) \otimes_R N \to H_n(A_\bullet \otimes_R N).$$

It's induced by the natural maps $(\ker \partial_n) \otimes_R N \to \ker(\partial_n \otimes 1)$ and $(\operatorname{im} \partial_{n+1}) \otimes_R N \to \operatorname{im}(\partial_{n+1} \otimes 1)$. If we think of the target spaces as submodules of $A_n \otimes_R N$, then those maps are induced by the respective inclusions $\ker \partial_n, \operatorname{im} \partial_{n+1} \hookrightarrow A_n$. Then the map above is given by

$$[a] \otimes m \mapsto [a \otimes m],$$

where a is an n-cycle of $A_\bullet$ and $[a]$ denotes its residue in homology. We can check directly that this map is well-defined: if $[a] = [b]$, then $a - b$ is a boundary in A_n, hence $[a \otimes m] - [b \otimes m] = [(a - b) \otimes m] = 0$ since $(a - b) \otimes m$ is a boundary in $A_n \otimes_R N$.

Part (iii) says that the above map is an isomorphism when $(-) \otimes_R N$ is exact (i.e. when N is **flat**).

Let's talk about exactness properties of (co)limits. Recall from Exercise 1.3.7 that limits commute. Since kernels are limits, it follows that *limits are left exact*. What does this mean? Let $\mathcal{C}$ be an abelian category with limits, and suppose $\mathcal{I}$ is an index category and A, B, C are diagrams in $\mathcal{C}$ indexed by $\mathcal{I}$. Let $f\colon A \to B$ and $g\colon B \to C$ be natural transformations, i.e. morphisms in $\mathcal{F}(\mathcal{I}, \mathcal{C})$. Explicitly, we have a sequence $A_i \to B_i \to C_i$ for each $i \in \mathcal{I}$, and these maps are compatible with the other maps in the diagrams. One can show, exactly as in Exercise 1.5.11, that $\mathcal{F}(\mathcal{I}, \mathcal{C})$ is an abelian category, and thus it makes sense to speak of exactness on the level of *diagrams*. In particular, suppose $f\colon A \to B$ is the kernel of $g\colon B \to C$. (On the object level, this means that f_i is the kernel of g_i for all $i \in \mathcal{I}$.) Then the assertion "limits are left exact" is the assertion that $\varprojlim A \to \varprojlim B$ is a kernel of $\varprojlim B \to \varprojlim C$.

How to remember this? Associate limits with kernels and kernels with left exactness. Dually, colimits are right exact; associate colimits with cokernels with right exactness.

In special cases (think: concrete categories?), filtered colimits are actually exact.

Exercise 1.5.16. Show that filtered colimits over $\mathbf{Mod}_A$ are exact.

Solution. By the discsussion above, colimits are right exact, so we need show left exactness. We'll use the explicit construction for colimits in $\mathbf{Mod}_A$ given in Exercises 1.3.4 and 1.3.5. Assume the notation from above. The map $f_*\colon \varinjlim A \to \varinjlim B$ sends $[(a_i, i)] \mapsto [(f_i(a_i), i)]$, and similarly for $g_*\colon \varinjlim B \to \varinjlim C$. We wish to show that $\ker(g_*)$, as a submodule of $\varinjlim B$, consists precisely of those elements $[(b_i, i)]$ in the image of f_*.

Suppose $[(g_i(b_i), i)] = 0 = [(0, j)]$. As $\mathcal{I}$ is filtered, there exists a map $i \to k$ such that $C_i \to C_k$ sends $g_i(b_i)$ to 0. Then the corresponding map $B_i \to B_k$ in B sends $b_i \mapsto b_k$ for some b_k, and the commuting square yields $g_k(b_k) = 0$.

$$\begin{array}{ccc} b_i & \longmapsto & g_i(b_i) \\ \downarrow & & \downarrow \\ b_k & \longmapsto & 0 \end{array}$$

In particular, as f_k is the kernel of g_k, we have $b_k \in \mathrm{im}(f_k)$ and thus

$$[(b_i, i)] = [(b_k, k)] \in \mathrm{im}(f_*),$$

as desired. □