HOMOLOGICAL ALGEBRA

ANDREW LEE

Abstract. In algebraic topology, homology groups are often used to encode information about a topological space, utilizing certain constructs called chain complexes, which capture certain types of information about the topological space. These chain complexes can then be used to form homology groups, which are some very interesting algebraic invariants of topological spaces, often conveying some useful information about the underlying topological space, such as connectedness.

Now, in the study of homological algebra, we study the homology of more general chain complexes, which encode information about more general algebraic structures, such as abelian groups, R-modules, or in general, any object of an abelian category.

In this paper, we lay out the fundamentals of modules, chain complexes, exact sequences, and homology, then build up to the functors Hom, Ext, Tensor, and Tor. The overall pedagogical approach to the topic was heavily inspired by [CS] and [Lin], and the other sources were used to round out the picture and find more examples.

1. R-Modules

In this section, we will give a quick review on the definition of a R-module. Note that unless otherwise specified, we will refer to commutative rings with unity as simply rings.

Definition 1.1. Let R be a ring. A *left R-module* M_R is an abelian group $(M, +)$ along with a scalar multiplication map $R \times M \to M$ such that $(r, m) \mapsto rm$. This scalar multiplication should also satisfy the following axioms, for $r, s \in R$ and $m, n \in M$.

(1) $r(m + n) = rm + rn$ (Left Distributivity)
(2) $(r + s)m = rm + sm$ (Right Distributivity)
(3) $(rs)m = r(sm)$ (Associativity)
(4) $1_R m = m$ (Identity)

A similar formulation can define right R-modules, but since we are dealing with commutative rings, these two formulations are equivalent. In this case, where R is a commutative ring, we simply refer to M_R as an R-module. Moreover, if the context is clear, we will often omit the R and call M the R-module.

Example. If $R = K$ is a field, these K-modules are precisely the vector spaces over K.

Example. We may also think of abelian groups as $\mathbb{Z}$-modules, with the scalar multiplication map being the (perhaps) obvious one, where $n \in \mathbb{Z}$ maps m to a sum of m with itself n times: $(n, m) \mapsto m + \cdots + m$. In fact, as we will see later with different formulations of homology groups, this is why R-modules are a natural generalization of abelian groups.

Just as with groups, R-modules have similar notions of R-submodules. These aren't particularly interesting; they are nearly identical to formulations over other algebraic objects.

Date: July 2020.

Definition 1.2. Let M_R be an R-module. An *R-submodule* is an abelian subgroup $N \subseteq M$ such that for all $r \in R$ and $n \in N$, $rn \in N$. That is, elements in N are closed under scalar multiplication by elements in the ring.

Example. Ideals of R can be considered R-submodules, with the scalar multiplication map simply being the multiplication on R.

Definition 1.3. Let M and N be R-modules. A map $\phi\colon M \to N$ is called an *R-linear map* or an *R-module homomorphism* if for all $m, m' \in M$ and $r \in R$:

(1) $\phi(m + m') = \phi(m) + \phi(m')$
(2) $\phi(rm) = r\phi(m)$

We denote by $\mathrm{Hom}_R(M, N)$ the set of all R-module homomorphisms from $M \to N$.

Note that we can naturally see a category R-**Mod** with objects R-modules and morphisms R-module homomorphisms.

One class of R-modules that is of interest is the *finitely generated R-modules*, which we will soon define.

Definition 1.4. Let $\{x_1, \ldots, x_m\}$ be a collection of elements from an R-module M. Consider the subset *generated by* these elements:

$$\langle x_1, \ldots, x_m \rangle = \left\{ \sum_{i=1}^{m} r_i x_i \mid r_i \in R \right\}$$

This subset is an R-submodule of M.

Moreover, if the subset $\langle x_1, \ldots, x_m \rangle$ is equal to the R-module M, we say that the set $\{x_1, \ldots, x_m\}$ *spans* or is a *spanning set* of M.

This leads us right into the definition of a finitely generated module.

Definition 1.5. Let M be an R-module. If there exists a finite spanning set $\{x_1, \ldots, x_m\}$ of M, then M is *finitely generated*.

Example. The Cartesian product $\mathbb{Z} \times \mathbb{Z}$ is a finitely generated $\mathbb{Z}$-module, and is spanned by the two elements $x_1 = (1, 0)$ and $x_2 = (0, 1)$.

Example. By the Fundamental Theorem of finitely generated abelian groups, an abelian group M is finitely generated if and only if it is of the form

$$M = \mathbb{Z}^n \times \mathbb{Z}/n_1\mathbb{Z} \times \cdots \times \mathbb{Z}/n_r\mathbb{Z}$$

with $n_i \geq 1$. Thus, any finitely generated abelian group contains a finite number of copies of $\mathbb{Z}$, along with a torsion part which is the product of a finite number of cyclic groups of finite order. Moreover, M is spanned by the set $x_i = (0, \ldots, 0, 1, 0, \ldots, 0)$ with a 1 in the ith position, then these elements span the module M.

Note that this aligns closely with the notions of spanning subsets of $\mathbb{R}^n$ from linear algebra.

Definition 1.6. Let $\{x_1, \ldots\}$ be the generating set of an R-submodule of M. If $\{x_1, \ldots\}$ spans M and is linearly independent over R, we say that $\{x_1, \ldots\}$ forms an *R-module basis* of M.

If an R-module M admits a basis, the module is called a *free module*.

Note that free modules need not be finitely generated, but there is a slight caveat: the elements of a free module can only be *formal* linear combinations of the basis elements; that is, all but finitely many elements in the generating set must have coefficient zero.

Example. The elements $x_1 = (1,0)$ and $x_2 = (0,1)$ form a $\mathbb{Z}$-module basis of $\mathbb{Z} \times \mathbb{Z}$. Thus, $\mathbb{Z} \times \mathbb{Z}$ is free.

Example. The zero module $M = 0$ is a free module, with basis the empty set.

Nonexample. An arbitrary finitely generated abelian group M of the form given by the Fundamental Theorem of finitely generated abelian groups, in general, does not always admit a basis. In fact, M only admits a basis if it has zero torsion part. (A consequence of Proposition 6.12.) As we will see later, the Tor functor is thus named because Tor_1 can extract information about the torsion part of an R-module, with R an integral domain.

2. Chain Complexes and Exact Sequences

Chain complexes are ubiquitous in homological algebra, since they are able to capture a vast wealth of information about certain objects. In algebraic topology, these chain complexes (as well as cochain complexes) are used to extract information about topological spaces in the form of certain abelian groups. For our purposes, we will introduce these concepts with chain complexes over R-modules, but these results are easily generalizable to any abelian category.

Definition 2.1. Let $(A_\bullet, d_\bullet)$ be a family of modules A_i and module homomorphisms $d_i \colon A_i \to A_{i-1}$ indexed by the integers. The module A_i called the *degree* i component of the family, and the homomorphism d_i is called the *boundary operator in degree* i, or simply the i*th boundary operator*.

Such a family $(A_\bullet, d_\bullet)$ is a *chain complex* if $d_{i-1} \circ d_i = 0$ for all i. Sometimes, we may refer to the complex simply as $A_\bullet$.

Alternatively, for the dual formulation $(A^\bullet, d^\bullet)$ with $d^i \colon A^i \to A^{i+1}$, $(A^\bullet, d^\bullet)$ is a *cochain complex* if $d_{i+1} \circ d_i = 0$ for all i. Ignoring indices, we can express this condition as $d^2 = 0$. Schematically, a chain complex can be displayed as follows.

$$\cdots \xleftarrow{d_0} A_0 \xleftarrow{d_1} A_1 \xleftarrow{d_2} A_2 \xleftarrow{d_3} \cdots$$

Most of the following discussion will not concern itself with the difference between chains and cochains, since that is simply a matter of distinguishing between increasing or decreasing boundary operators.

The distinction becomes more important when definining homology or cohomology theories, which will come later. The important defining property of complexes is the identity $d^2 = 0$. However, there is another equivalent formulation, shown below.

Remark 2.2. Let $f \colon M \to N$ and $g \colon N \to P$ be two module homomorphisms. The composition $g \circ f$ is the zero map if and only if $\mathrm{im}(f) \subseteq \ker(g)$. Moreover, in the context of (co)chain complexes, the composition of two boundary maps is zero if and only if the image of the first is contained in the kernel of the next.

This statement is rather simple to prove, so we will not do it here. However, this new view of the $d^2 = 0$ condition motivates new questions. Given a certain chain complex, when is the image of one boundary map equal to the kernel of the next? If they are not equal, how much do they "fail" to be equal? These are the central motivating questions behind our use of homology.

Let us first recall some notation and definitions.

Definition 2.3. Let $\phi\colon M \to N$ be an R-module homomorphism. Then, define the *kernel*, *image*, and *cokernel* as follows.

(1) $\ker(\phi) = \{m \in M \mid \phi(m) = 0\}$
(2) $\operatorname{im}(\phi) = \{\phi(m) \in N \mid m \in M\}$
(3) $\operatorname{coker}(\phi) = N/\operatorname{im}(\phi)$

Recall that $\ker(\phi) = \{0\}$ if and only if ϕ is an injective homomorphism. Moreover, ϕ is surjective if and only if $\operatorname{im}(\phi) = N$; that is, if $\operatorname{coker}(\phi) = \{0\}$. In this way, the kernel and cokernel can be seen as duals. We will soon see that the kernel and cokernel play a special role when examining sequences.

Definition 2.4. A complex $(A_\bullet, d_\bullet)$ is said to be *exact* at A_i if $\operatorname{im}(d_{i+1}) = \ker(d_i)$. The complex is called an *exact sequence* if it is exact for all i.

Remark 2.5. Exact sequences are complexes, but complexes are not necessarily exact. This "failure" of certain complexes to be exact is the central study of homology.

Exact sequences can lead to nice properties on the objects and boundary maps. First, let us consider some simple cases where the sequence is finite and begins or ends with the zero module.

Example. Given R-modules A, B, and C, the following are true. (Note that these need not be R-modules, and can be objects of any abelian category.)

(1) $0 \longrightarrow A \xrightarrow{\alpha} B$ is exact if and only if α is injective.
(2) $B \xrightarrow{\beta} C \longrightarrow 0$ is exact if and only if β is surjective.
(3) $0 \longrightarrow A \xrightarrow{\phi} B \longrightarrow 0$ is exact if and only if ϕ is an isomorphism.

Proof. $(1, \Rightarrow)$ Suppose the sequence is exact. Then, the image of the inclusion $\iota\colon 0 \to A$ is $\operatorname{im}(\iota) = \{0\}$, and thus by exactness, $\ker(\alpha) = \{0\}$ and α is injective.

$(1, \Leftarrow)$ Now suppose α is injective. Then, $\ker(\alpha) = \{0\}$ and the image of the inclusion is also trivially $\{0\}$. Thus they are equal and the sequence is exact.

The proof of (2) is similar to that of (1), and $(1), (2) \implies (3)$ is trivial. □

Note that if we have two consecutive nonzero objects in a chain, mapping into/out of zero on both sides, the two objects are necessarily isomorphic. Moreover, we get an interesting outcome with three consecutive nonzero objects, interesting enough that we have a name for it.

Definition 2.6. A *short exact sequence* is an exact sequence of the form

$$0 \longrightarrow A \xrightarrow{\alpha} B \xrightarrow{\beta} C \longrightarrow 0$$

By the definition of exactness and the examples above, α is injective, β is surjective, and $\operatorname{im}(\alpha) = \ker(\beta)$. It can be shown that β induces a natural quotient structure on C, since β is surjective and therefore $B/\ker(\beta) \cong \operatorname{im}(\beta) \cong C$. Using exactness, we may rewrite this as $C \cong B/\operatorname{im}(\alpha)$. Intuitively, we can see A as a partition of B, with the inclusion map α, then a map β which induces a quotient structure on C.

Conversely, if $C \cong B/A$, there is a short exact sequence

$$0 \longrightarrow A \hookrightarrow B \twoheadrightarrow C \longrightarrow 0$$

where $\hookrightarrow$ shows an injective map, and $\twoheadrightarrow$ shows a surjective map.

Example. Given a sequence, we can often construct induced exact sequences based on the objects and maps. For example, consider $\phi\colon M \to N$, with ϕ an R-module homomorphism. Then, the following sequence is exact.

$$0 \longrightarrow \ker(\phi) \longrightarrow M \xrightarrow{\phi} N \longrightarrow \operatorname{coker}(\phi) \longrightarrow 0$$

It follows that the case above, where $0 \longrightarrow M \xrightarrow{\phi} N \longrightarrow 0$ is exact, corresponds exactly to when $\ker(\phi) = \operatorname{coker}(\phi) = \{0\}$.

Definition 2.7. A *long exact sequence* is an exact sequence containing more than three (often infinitely many) nonzero terms, and can be written as follows.

$$\cdots \xrightarrow{d_0} A_1 \xrightarrow{d_1} A_2 \xrightarrow{d_2} \cdots$$

These long exact sequences come up quite often and are rather useful. Interestingly, they can also be seen as a sequence of short exact sequences.

Example. Let $(A_\bullet, d_\bullet)$ be a long exact sequence, and let the boundary maps be increasing: $d_i\colon A_i \to d_{i+1}$. Define $K_i = \ker(d_i) = \operatorname{im}(d_{i-1})$. Then, the long exact sequence induces a sequence of short exact sequences in the diagram below, where each diagonal is a short exact sequence.

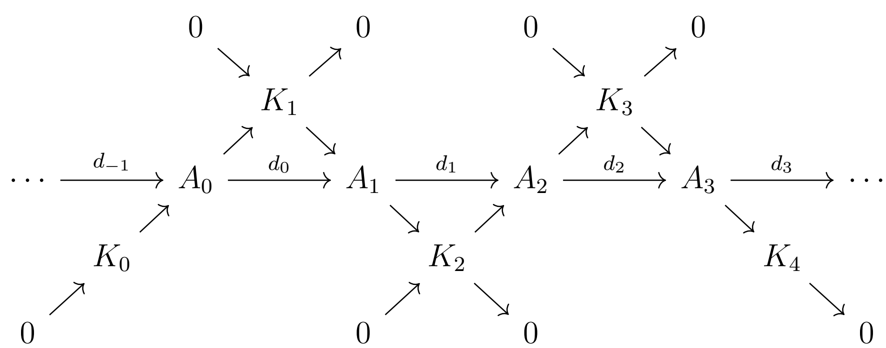

Conversely, a sequence of short exact sequences can be combined into a long exact sequence by taking their middle objects.

Another way of making a long exact sequence is given by the Snake Lemma.

Lemma 2.8 (Snake Lemma)**.** *Let the following diagram be a commutative diagram over the category R-**Mod** (or in general some abelian category), such that the rows are exact.*

$$\begin{array}{ccccccccc} & & A & \xrightarrow{f} & B & \xrightarrow{g} & C & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle\alpha} & & \downarrow{\scriptstyle\beta} & & \downarrow{\scriptstyle\gamma} & & \\ 0 & \longrightarrow & A' & \xrightarrow{f'} & B' & \xrightarrow{g'} & C' & & \end{array}$$

Then, there exists a long exact sequence of their kernels and cokernels with a map $s\colon \ker(\gamma) \to \operatorname{coker}(\alpha)$ (shown below in dark blue) such that the following diagram commutes, and the *sequence in blue is exact. Moreover, the blue sequence along with the red arrow (equiv. the green arrow) is exact if and only if f is injective (equiv. if g' is surjective).*

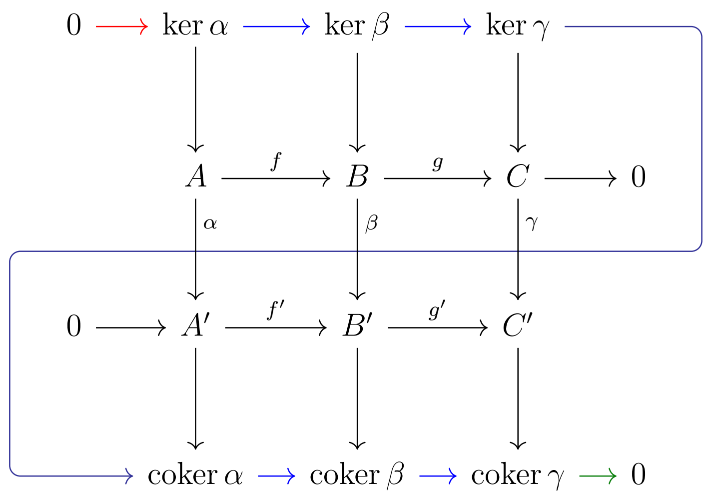

Proof. (Kate Gunzinger, *It's My Turn* (1980), [Gun])[1]

Dr. Gunzinger: "Let me just show you how to construct the map s, which is the fun of the lemma anyhow, okay? So you assume you have an element in $\ker(\gamma)$—that is, an element $c \in C$ such that γ takes you to 0 in C'. You pull it back to B by the map g, which is surjective, with $b \in B$ uniquely determined up to addition by an element in $\operatorname{im}(f)$. Alright? So we pull it back to a fixed $b \in B$ here, then you take $\beta(b)$, which takes you to 0 in C' by the commutivity of the diagram! It's therefore in $\ker(g')$, hence it's in $\operatorname{im}(f')$ by the exactness of the lower sequence, so we can pull it back to an element in A'—"

Cooperman: "No...no—it's *not well defined*."

Gunzinger: "which it turns out is well defined up to the elements of $\operatorname{im}(\alpha)$. And thus defines an element in $\operatorname{coker}(\alpha)$, and that's the snake!"

Author: "This is not part of the movie, but the red and green conditions are routine to check, and come down to some diagram chasing."

□

3. Homology

Earlier, we eluded to the fact that homology somehow measures how much certain objects in a complex fail to be exact. In this section, we will formalize that notion.

Definition 3.1. Let $(A_\bullet, d_\bullet)$ be a complex with descending boundary operators ($d_i : A_i \to A_{i-1}$). We define the R-submodule of *n-cycles* to be $Z_n(A_\bullet) = \ker(d_n)$ and the R-submodule of *n-boundaries* to be $B_n(A_\bullet) = \operatorname{im}(d_{n+1})$.

Since this is a complex, we have $B_n(A_\bullet) \subseteq Z_n(A_\bullet)$ for all n. Therefore, we may describe their quotient; this is precisely the homology.

Definition 3.2. Let $(A_\bullet, d_\bullet)$ be a complex. Then, the *nth homology module* of the complex is defined as the quotient R-module

$$H_n(A_\bullet) := Z_n(A_\bullet)/B_n(A_\bullet) = \ker(d_n)/\operatorname{im}(d_{n+1})$$

[1]The original idea to include this proof of the snake lemma is from [Cla].

Remark 3.3. Note that if $Z_n = B_n$, the homology module is zero. Thus, $A_\bullet$ is an exact sequence if and only if $H_n(A_\bullet) = 0$ for all n.[2]

We see that these homology modules can give us a measure of how close a certain complex may come to being exact. Intuitively, this is similar to how a commutator subgroup can tell us how close a certain group is to being abelian; if the commutator is zero, then every pair of elements commutes and the group is abelian.

In fact, many (co)homology theories are based on the same idea of measuring the "failure to be exact" of certain (co)chain complexes that have some sort of interpretation in the theory.

For example, the nth simplicial homology group of some (triangularizable) topological space X is a measure of how many n-dimensional "holes" are in the space. In this case, exactness at A_n is interpreted as a topological space having no n-dimensional holes.

Another example is de Rham cohomology, which (roughly speaking) measures the extent to which the fundamental theorem of calculus fails in higher dimensions. Again, here the notion of exactness has another interpretation, related to the fundamental theorem of calculus.

In the remainder of this text, we will not concern ourselves with the geometric or topological side of homology, instead focusing on the information we can extract from complexes using these quotient homology modules.

4. The Hom Functor

Recall that from basic category theory, the morphisms between objects A and B of a category $\mathscr{C}$ live in the set $\mathrm{Hom}_{\mathscr{C}}(A, B)$. If we fix A, then we get a (covariant) Hom functor, $\mathrm{Hom}_{\mathscr{C}}(A, -)$, and dually if we fix B, we get the contravariant Hom functor, $\mathrm{Hom}_{\mathscr{C}}(-, B)$. In general, we can define the Hom bifunctor $\mathrm{Hom}_{\mathscr{C}}(-, -)$, which is contravariant in one argument and covariant in the other, defined on a product category with $\mathrm{Hom}_{\mathscr{C}}(-, -)\colon \mathscr{C}^{\mathrm{op}} \times \mathscr{C} \to \mathbf{Set}$. However, in this text, we will only consider the (covariant/contravariant) Hom functors, rather than the bifunctor.

Let us review the properties of the covariant Hom functor. Let us fix an R-module M. Then, for any R-module N, the set of R-module homomorphisms from M to N, $\mathrm{Hom}_R(M, N)$, forms an R-module under pointwise operations. Moreover, given an R-module homomorphism $\phi\colon N_1 \to N_2$, there is an induced R-module homomorphism $\phi_*\colon \mathrm{Hom}_R(M, N_1) \to \mathrm{Hom}_R(M, N_2)$ such that the following square commutes.

$$\begin{array}{ccc} N_1 & \xrightarrow{\quad\phi\quad} & N_2 \\ \big\downarrow {\scriptstyle \mathrm{Hom}_R(M,-)} & & \big\downarrow {\scriptstyle \mathrm{Hom}_R(M,-)} \\ \mathrm{Hom}_R(M, N_1) & \xrightarrow{\phi_*} & \mathrm{Hom}_R(M, N_2) \end{array}$$

Specifically, a homomorphism $\alpha\colon M \to N_1$ gets mapped to a homomorphism $\phi_*(\alpha) = \phi \circ \alpha$.

Dually, the contravariant Hom fuctor will send an R-module homomorphism $\phi\colon M_1 \to M_2$ to a homomorphism in the opposite direction; that is, $\phi^*\colon \mathrm{Hom}(M_2, N) \to \mathrm{Hom}(M_1, N)$, again given by composition: $\phi^*(\alpha) = \alpha \circ \phi$.[3]

[2]This will become a recurring theme in our study of Tor and Ext, which themselves are defined as the homology of certain sequences.

[3]Note that in the covariant case, the star is a subscript, and in the contravariant case, the star is in the superscript. This general convention will be followed throughout this text.

Now that we have seen the action of the Hom functor on R-modules and homomorphisms between them, let us consider their action on exact sequences, and just how much of that "exactness" is preserved.

Proposition 4.1. *Let M, N be R-modules, and let*

$$0 \longrightarrow A \xrightarrow{\alpha} B \xrightarrow{\beta} C \longrightarrow 0$$

be a short exact sequence of R-modules. Then, we get an exact sequence through the covariant Hom functor.

$$0 \longrightarrow \mathrm{Hom}_R(M, A) \xrightarrow{\alpha_*} \mathrm{Hom}_R(M, B) \xrightarrow{\beta_*} \mathrm{Hom}_R(M, C)$$

We call such a functor a (covariant) left exact functor*. Moreover, if we fix N and consider the contravariant Hom functor, we again get an exact sequence with some arrows reversed.*

$$0 \longrightarrow \mathrm{Hom}_R(C, N) \xrightarrow{\beta^*} \mathrm{Hom}_R(B, N) \xrightarrow{\alpha^*} \mathrm{Hom}_R(A, N)$$

This functor is also said to be a (contravariant) left exact functor*.*

This proposition can be proved fairly easily, simply requiring us to check exactness at $\mathrm{Hom}_R(M, B)$ (or $\mathrm{Hom}_R(B, N)$), and injectivity of α_* (or β^*).

Dual to the notion of a left exact functor is that of a *right exact functor*, which sends an exact sequence $0 \longrightarrow A \longrightarrow B \longrightarrow C \longrightarrow 0$ to $\mathscr{F}(A) \longrightarrow \mathscr{F}(B) \longrightarrow \mathscr{F}(C) \longrightarrow 0$, for some functor $\mathscr{F}$. Of course, if a functor is both left and right exact, we say that it is an *exact functor*.

5. Resolutions and Ext

In this section, we aim to introduce some different types of resolutions (free, projective, and injective), and their relation to the Ext functor.

Let us begin with free resolutions, to introduce the concept of a resolution.

Definition 5.1. Let M be an R-module. If there is a complex $(F_\bullet, d_\bullet)$ (with decreasing boundary operators) such that the following sequence is exact, then the complex $F_\bullet$ (with indices $n \geq 0$, and not including M) is said to be a *resolution* of M.

$$\cdots \xrightarrow{d_3} F_2 \xrightarrow{d_2} F_1 \xrightarrow{d_1} F_0 \xrightarrow{\varepsilon} M \longrightarrow 0$$

The homomorphism ε is called the augmentation map, and is defined as the projection $\varepsilon \colon F_0 \to \mathrm{coker}(d_1)$. Succinctly, we can denote a resolution as follows.

$$F_\bullet \xrightarrow{\varepsilon} M \longrightarrow 0$$

Furthermore, if all the R-modules F_i are free, then we say that $F_\bullet$ is a *free resolution* of M.

Remark 5.2. Note that the resolution does not actually contain M, or for that matter the map ε. (In fact, sometimes we may refer to $\varepsilon = d_0$; the convention to use ε is to emphasize that ε is not in the resolution.) However, we may recover M as the cokernel $M = \mathrm{coker}(d_1)$. This follows from exactness. Since $F_0 \xrightarrow{\varepsilon} M \to 0$ is exact, it follows that ε is surjective. Therefore, $M \cong \mathrm{im}(\varepsilon) \cong F_0/\ker(\varepsilon) \cong F_0/\mathrm{im}(d_1) \cong \mathrm{coker}(d_1)$.

Now, which kinds of modules admit free resolutions? As it turns out, all modules admit free resolutions. Before we prove this, though, we must prove a quick lemma.

Lemma 5.3. *Every module is a quotient of a free module.*

Proof. We wish to show that any module M is the quotient of a free module F. Let $X = M \setminus \{0\}$ be the set of nonzero elements in M. Then, let $F = \langle X \rangle$ be the free module generated by X. These structures induce two canonical maps: the inclusion $X \overset{\iota}{\hookrightarrow} M$ and the projection $F \overset{f}{\twoheadrightarrow} M$ with

$$\iota\colon \tilde{x}_i \mapsto x_i$$

where $\tilde{x}_i$ denotes the (nonzero) element $x_i \in M$ restricted to X, and f the projection map below.

$$f\colon \sum r_i \tilde{x}_i \mapsto \sum r_i x_i$$

Clearly, f is surjective, and thus its image is equal to M. Therefore, there is a natural quotient

$$M \cong F/\ker(f)$$

and thus any module is the quotient of a free module. □

With this, we may prove the following proposition.

Proposition 5.4. *Any module over any ring admits a free resolution.*

Proof. By the above lemma, M may be written as the quotient of a free module F_0 by the kernel of the homomorphism $\varepsilon\colon F_0 \to M$; that is, $M \cong F_0/K_0$. Since the kernel $K_0 = \ker(\varepsilon)$ is itself an R-module, we can repeat the process to K_0 and write it as the quotient of a free module F_1 by the kernel $\ker(\pi_1)$.

Now consider the inclusion map $\iota_1\colon K_0 \to F_0$, and the induced composition map $d_1 = \iota_1 \circ \pi_1\colon F_1 \to F_0$. Since ι_1 is injective, it follows that $\ker(d_1) = \{x \in F_1 \mid \iota_1(\pi_1(x)) = 0\} = \{x \in F_1 \mid \pi_1(x) \in \ker(\iota_1)\} = \{x \in F_1 \mid \pi_1(x) = 0\} = \ker(\pi_1)\}$. We call this new module $K_1 = \ker(d_1) = \ker(\pi_1)$.

From here, it follows that we can repeat this process inductively, resulting in the complex shown below.

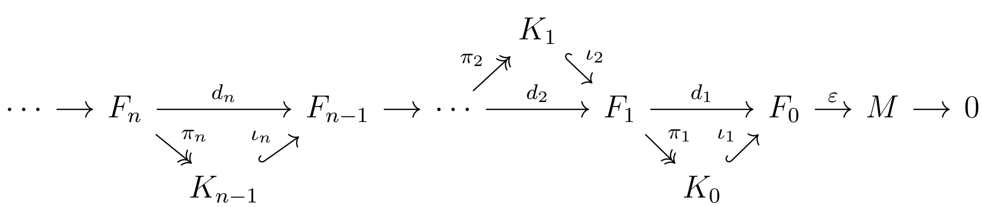

In particular, we define K_{n-1} as the kernel $K_{n-1} = \ker(d_{n-1})$ (or $\ker(\pi_{n-1})$; they are equivalent since the $\iota_\bullet$ are injective), then let F_n be the free module such that $F_n/\ker(\pi_n) \cong K_{n-1}$, with $\pi_n\colon F_n \to K_{n-1}$ the canonical quotient map, which is surjective. Then, define $\iota_n\colon K_{n-1} \to F_{n-1}$ be the canonical inclusion map, which is injective.

Now notice that through this process, we have a "zigzag" of surjective and injective maps. Recall that $0 \longrightarrow A \xrightarrow{\alpha} B$ is exact if and only if α is injective, and $B \xrightarrow{\beta} C \longrightarrow 0$ is exact if and only if β is surjective. Therefore, we may express the above complex as a sequence of short exact sequences, as shown below. Thus, the entire sequence is exact and therefore M

admits a free resolution.

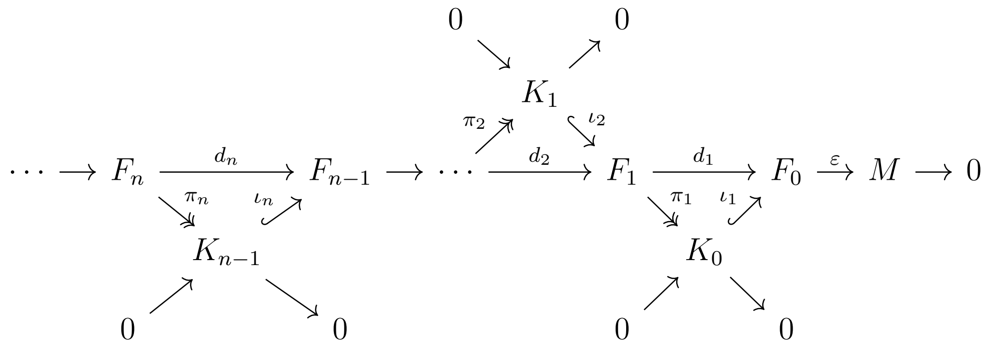

∎

Now let us consider projective resolutions.

Definition 5.5. An R-module P is *projective* if for any surjection of R-modules $\pi\colon Q' \twoheadrightarrow Q$, any homomorphism $\phi\colon P \to Q$ induces a homomorphism $\phi'\colon P \to Q'$ such that the diagram below commutes.

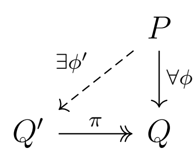

It is worth mentioning that this is not a universal property; ϕ' is not necessarily unique.

Example. A free module F satisfies the condition of being projective.

Proof. Suppose F is a free R-module with basis $\langle x_1, \ldots, x_m \rangle$. Then, let $\pi\colon Q' \twoheadrightarrow Q$ be a surjection. For some arbitrary basis element x_i, consider $\phi(x_i) \in Q$. Since π is surjective, there exists at least one $y_i \in Q'$ such that $\pi(y_i) = \phi(x_i)$. Now define $\phi'\colon x_i \mapsto y_i$. Then, $\pi(\phi'(x_i)) = \pi(y_i) = \phi(x_i)$, and thus $\pi \circ \phi' = \phi$ and the diagram commutes.

Therefore, any free module is also projective. ∎

Now, recall our exact functors from earlier. Recall that the Hom functor (both variants) is left exact, but not necessarily right exact. Of course, this begs the question: when is the Hom functor right exact (and therefore exact)?

Theorem 5.6. *An R-module P is projective if and only if the covariant functor* $\mathrm{Hom}_R(P, -)$ *is exact.*

Proof. We begin with the forward direction. Let P be a projective R-module. Of course, it suffices to simply prove that $\mathrm{Hom}_R(P, -)$ is right exact given that it is left exact. Therefore, it suffices to prove that given the exact sequence $0 \longrightarrow A \longrightarrow B \longrightarrow C \longrightarrow 0$, the following sequence is exact at $\mathrm{Hom}_R(P, C)$.

$$0 \longrightarrow \mathrm{Hom}_R(P, A) \xrightarrow{\alpha_*} \mathrm{Hom}_R(P, B) \xrightarrow{\beta_*} \mathrm{Hom}_R(P, C) \longrightarrow 0$$

Thus, it suffices to show that β_* is surjective, or that $\mathrm{Hom}_R(P, -)$ preserves surjections.

Let $\pi\colon A \twoheadrightarrow B$ be a surjection on R-modules. Then, we want to show that $\pi_*\colon \mathrm{Hom}_R(P, A) \to \mathrm{Hom}_R(P, B)$ is also a surjection. Let $\phi\colon P \to B$ be a homomorphism in $\mathrm{Hom}_R(P, B)$, giving us the diagram below. Then, since P is surjective, there exists some map $\phi'\colon P \to A$ such that the diagram commutes.

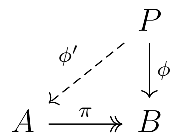

In this situation, $\pi_*(\phi') = \phi$, and thus π_* is a surjection.

Now, let us consider the other direction. Suppose $\mathrm{Hom}_R(P, -)$ is exact. Then, it preserves surjections; $\pi\colon A \twoheadrightarrow B$ maps to a surjection $\pi_*\colon \mathrm{Hom}_R(P, A) \twoheadrightarrow \mathrm{Hom}_R(P, B)$ via our functor. Then, since π_* is surjective, for every $\phi \in \mathrm{Hom}_R(P, B)$, there exists another $\phi' \in \mathrm{Hom}_R(P, A)$ such that $\phi = \pi_*(\phi') = \pi \circ \phi'$, which makes the same above diagram commute.

Thus P is projective if and only if $\mathrm{Hom}_R(P, -)$ is exact. □

Definition 5.7. Let M be an R-module, and $P_\bullet$ a resolution of M:

$$P_\bullet \xrightarrow{\varepsilon} M \longrightarrow 0$$

If all the R-modules P_i are projective, then we say that $P_\bullet$ is a *projective resolution* of M.

Proposition 5.8. *Any module over any ring admits a projective resolution.*

Proof. Given that any module over any ring admits a free resolution, and all free modules are projective, this result follows immediately. □

As a quick side note, one can define the projective dimension of a module, denoted $\mathrm{pd}(M)$, as the shortest length of any projective resolution of M, which turns out to have vast connections to other notions of dimension in algebraic geometry. Unfortunately, though, that is outside the scope of this paper.

Given that any module M admits a projective resolution, we can consider what happens when we apply the Hom functor to our projective resolution.

Definition 5.9. Let M be an R-module, and let $P_\bullet$ be a projective resolution of M.

Fix an R-module N, and consider the complex induced by the contravariant Hom functor $\mathrm{Hom}_R(-, N)$.

$$0 \longrightarrow \mathrm{Hom}_R(P_0, N) \xrightarrow{d_1^*} \mathrm{Hom}_R(P_1, N) \xrightarrow{d_2^*} \cdots$$

Then, define the homology groups $\mathrm{Ext}_R^n(M, N)$, for $n \geq 0$, as follows.

$$\mathrm{Ext}_R^n(M, N) := \ker(d_{n+1}^*)/\mathrm{im}(d_n^*)$$

Proposition 5.10. *The homology groups* $\mathrm{Ext}_R^n(M, N)$ *are independent of the choice of projective resolution.*

Proof. The proof requires some technical machinery (chain homotopies) that was not presented in this text, and thus will not be presented here. □

Example. Given R-modules M and N, $\mathrm{Ext}_R^0(M, N) \cong \mathrm{Hom}_R(M, N)$.

Now, there is an alternate way to present the Ext functor, using injective resolutions instead of projective resolutions. We begin by discussing injective modules.

Definition 5.11. An R-module I is *injective* if for any injection of R-modules $\iota\colon Q \hookrightarrow Q'$, any homomorphism $\phi\colon Q \to I$ induces a homomorphism $\phi'\colon Q' \to I$ such that the diagram below commutes.

$$\begin{array}{ccc} I & & \\ {\scriptstyle\forall\phi}\uparrow & \nwarrow{\scriptstyle\exists\phi'} & \\ Q & \xhookrightarrow{\iota} & Q' \end{array}$$

This is the dual notion of a projective module, so we expect many of the results true of projective modules to be true of injective modules. Of course, we have a dual result to Theorem 5.6 above, which shows that an R-module P is projective if and only if $\mathrm{Hom}_R(P, -)$ is exact. The proofs of these results will be omitted from this text, as they are quite similar to the proofs for the projective version.

Note however that although the proof for Proposition 5.13 does not follow from the fact that all modules admit free resolutions (Proposition 5.4), the construction is similar, since any module can be embedded into an injective module. (Just like how any module is the quotient of a free module.)

Theorem 5.12. *An R-module I is injective if and only if the contravariant functor* $\mathrm{Hom}_R(-, I)$ *is exact.*

Proposition 5.13. *Any module over any ring admits an injective resolution.*

One large motivating reason to define Ext in terms of projective modules (rather than free modules), is that there is a natural dual.

Definition 5.14. Let N be an R-module, and $I_\bullet$ a resolution of N:

$$0 \longrightarrow N \xrightarrow{\varepsilon} I_\bullet$$

If all the R-modules I_i are injective, then we say that $I_\bullet$ is an *injective resolution* of N.

Again, we can use the Hom functor to derive a second formulation of the Ext functor.

Definition 5.15. Let N be an R-module, and let $I_\bullet$ be an injective resolution of N. Fix an R-module M, and consider the complex induced by the covariant Hom functor $\mathrm{Hom}_R(M, -)$.

$$0 \longrightarrow \mathrm{Hom}_R(M, I_0) \xrightarrow{d_{0,*}} \mathrm{Hom}_R(M, I_1) \xrightarrow{d_{1,*}} \cdots$$

Then, define the homology groups $\mathrm{Ext}_n^R(M, N)$, for $n \geq 0$, as follows. Note the difference from the previous definition; this time, the n is in the subscript.

$$\mathrm{Ext}_n^R(M, N) := \ker(d_{n,*})/\mathrm{im}(d_{n-1,*})$$

Moreover, these homology groups are independent of the choice of injective resolution.

Theorem 5.16. *For any R-modules M and N, and any integer $n \geq 0$, the two formulations of* Ext *are equal:*

$$\mathrm{Ext}_n^R(M, N) = \mathrm{Ext}_R^n(M, N)$$

Proof. The proof of this theorem is quite involved, and uses a fairly nasty diagram chase to achieve its goals. For a proof, see pg. 42 of [CS]. □

Now, at this point, a sensible reader might inquire: well, it's quite neat that you can define this homology using both projective and injective resolutions and get the same result, but why is it called Ext? Glad you asked.

It turns out that the Ext functor gets its name from its relation to extensions of modules.

Definition 5.17. Let M and N be R-modules. Then, an *extension of M by N* is a short exact sequence of modules

$$0 \longrightarrow N \longrightarrow E \longrightarrow M \longrightarrow 0$$

where two extensions are considered equivalent if the following diagram commutes.

$$\begin{array}{ccccccccc} 0 & \longrightarrow & N & \longrightarrow & E & \longrightarrow & M & \longrightarrow & 0 \\ & & \| & & \downarrow & & \| & & \\ 0 & \longrightarrow & N & \longrightarrow & E' & \longrightarrow & M & \longrightarrow & 0 \end{array}$$

Moreover, define $\operatorname{Ext}(M, N)$ to be the equivalence class of extensions of M by N.

Remark 5.18. Note that the Five Lemma (which we have not (yet) presented in this text) implies that the middle map is an isomorphism.

Proposition 5.19. *For R-modules M and N,* $\operatorname{Ext}(M, N)$ *and* $\operatorname{Ext}^1_R(M, N)$ *are naturally isomorphic.*

Therefore, the first Ext module has some interpretation as the equivalence classes of extensions of M by N. Now, a particularly inquisitive reader might now be wondering: well, what about the other Ext functors?

Unfortunately, those are a little more difficult. But, we do have some nice results when certain Ext groups are zero.

Proposition 5.20. *Let M and N be R-modules. If M is projective* or *N is injective, then for all $n > 0$,*

$$\operatorname{Ext}^n_R(M, N) = 0$$

In lieu of a formal proof, we will present some discussion to solidify the intuition behind this proposition. We will explain the result for M projective, since the argument for when N is injective is basically identical.

Recall that projective resolutions are exact by definition, and that $\operatorname{Hom}_R(M, -)$ is exact if and only if M is projective. Then, if our initial module M is projective, it follows that the Hom functor is exact and therefore gives us an exact sequence. Of course, by definition, the homology groups of an exact sequence are all zero, so clearly the Ext functor is zero at each object.

The converse of this is also true.

Proposition 5.21. *Fix R-modules M and N.*

(1) *If* $\operatorname{Ext}^1_R(M, N) = 0$ *for all N, then M is projective.*
(2) *If* $\operatorname{Ext}^1_R(M, N) = 0$ *for all M, then N is injective.*

Moreover, from Proposition 5.20, it follows that if $\operatorname{Ext}^1_R(M, N) = 0$ *for all M or all N, then* $\operatorname{Ext}^i_R(M, N) = 0$ *for all $i > 0$.*

Therefore, in some sense, we can see the Ext groups as measuring a sort of failure of M to be projective, where if M is projective, the sequence induced by Hom on a projective resolution is exact. Of course, we also have the dual case about whether N is injective.

In fact, in homological dimension theory, one can infer information about the homological "dimension" of a module by considering the least n such that $\operatorname{Ext}^i_R = 0$ for all $i > n$, and equate that to notions of projective or injective dimension.

6. Tensor Products and Tor

In this section, we begin by introducing the tensor product, then discussing its properties as a functor. Then, we will discuss the Tor functor, which is a sort of natural dual to the Ext functor.

This construction of the tensor product of two modules is due to [CS], which does a good job at making the tensor product intuitive.

Definition 6.1. Let M and N be two R-modules. Let $T(M,N)$ denote the free module indexed by elements in the product $M \times N$; that is,

$$T(M,N) := \bigoplus_{(x,y)\in M\times N} Re_{x,y}$$

where $e_{x,y}$ is simply denoting the element in $T(M,N)$ corresponding to the element $(x,y) \in M \times N$. Of course, this is an absolutely enormous module! Now consider the relation $\sim$ such that the following are true for all $x \in M$, $y \in N$, and $r \in R$:

$$\begin{aligned} e_{x+x',y} &\sim e_{x,y} + e_{x',y} \\ e_{x,y+y'} &\sim e_{x,y} + e_{x,y'} \\ re_{x,y} &\sim e_{rx,y} \sim e_{x,ry} \end{aligned}$$

Then, the *tensor product* of M and N is the quotient of $T(M,N)$ by the R-submodule induced by the relation $\sim$:

$$M \otimes_R N := T(M,N)/\sim$$

Of course, this relation ensures that the tensor product is distributive and *bilinear*; that is, it is linear in both terms.

Moreover, the equivalence class of the element $e_{x,y}$ is denoted $x \otimes y$, and is referred to as a *pure tensor*. In general, a *tensor* $z \in M \otimes_R N$ is the finite linear combination of pure tensors:

$$z = \sum_{i=1}^{n} x_i \otimes y_i$$

with $x_i \in M$ and $y_i \in N$. Thus the pure tensors generate the tensors in $M \otimes_R N$.

Just as Hom turns out to have functorial properties, the tensor product of modules over a ring R can be seen as a bifunctor $-\otimes_R-$. Moreover, there is a notion of duality between Hom and Tensor, which in turn motivates the dual notion of Ext but using the tensor product, Tor.

Remark 6.2. In the language of abstract nonsense (otherwise known as category theory), the relationship between Hom and Tensor can be described as a Tensor-Hom adjunction. Specifically, Tensor is the left adjoint to Hom, and Hom is the right adjoint to Tensor.

To take a mild detour further into the land of mathematical gibberish, the Hom functor commutes with arbitary limits, and Tensor commutes with arbitrary colimits. However, Hom does not (in general) commute with even finite colimits, and vice versa for Tensor. This failure to preserve short exact sequences motivates the construction of Ext and Tor, which measure the failure of these functors to preserve certain exact sequences (resolutions).

Coming back to the language we developed earlier in this text, we can show that the tensor product is a right exact functor.

Lemma 6.3. *Let $f\colon M \to M'$ and $g\colon N \to N'$ be two R-module homomorphisms. Then, we have an induced R-module homomorphism*

$$f \otimes g\colon M \otimes_R N \to M' \otimes_R N'$$

which sends $x \otimes y \mapsto f(x) \otimes g(y)$.

Proposition 6.4. *Let M, N be R-modules, and let*

$$0 \longrightarrow A \xrightarrow{\alpha} B \xrightarrow{\beta} C \longrightarrow 0$$

be a short exact sequence of R-modules. Then, the following induced sequences are also exact:

$$M \otimes_R A \xrightarrow{1\otimes\alpha} M \otimes_R B \xrightarrow{1\otimes\beta} M \otimes_R C \longrightarrow 0$$

$$A \otimes_R N \xrightarrow{\alpha\otimes 1} B \otimes_R N \xrightarrow{\beta\otimes 1} C \otimes_R N \longrightarrow 0$$

In other words, $M \otimes_R -$ and $- \otimes_R N$ are right exact covariant functors.

In fact, due to the commutativity of R, $M \otimes_R N = N \otimes_R M$, and thus these two constructions are equivalent.[4]

Of course, now that we have seen that Tensor is right exact, this begs the question: when is it exact? Well, thankfully, mathematicians came equipped to answer that question with—you guessed it—another definition.

Definition 6.5. Let N be an R-module, and let the following sequence be exact.

$$0 \longrightarrow A \longrightarrow B \longrightarrow C \longrightarrow 0$$

Now consider the induced sequence under the functor $- \otimes_R N$.

$$0 \longrightarrow A \otimes_R N \longrightarrow B \otimes_R N \longrightarrow C \otimes_R N \longrightarrow 0$$

If the induced sequence is exact, we call N a *flat* module.

Proposition 6.6. *Any projective module is flat.*

The proof of this proposition requires a little more machinery in the realm of split sequences, so we do not prove it here. However, we now have a little more information about different types of modules that we have encountered in this paper, namely, that free $\implies$ projective $\implies$ flat.

Recall that by Proposition 5.20, we can specify when $\mathrm{Ext}^i_R(M, N) = 0$ based on whether M is projective (or N is injective). This time, the dual notion is extremely simple to show. Let us define the dual notion to Ext, the homology groups Tor.

[4] In the remainder of this text, we shall use whatever convention that makes the most amusing word when the object is read out loud (i.e., corn).

Definition 6.7. Let K be a fixed R-module.[5] For any R-module N, let $P_\bullet$ be the projective resolution of N:

$$P_\bullet \xrightarrow{\varepsilon} N \longrightarrow 0$$

Applying the tensor product $- \oplus_R K$, we get the complex

$$\cdots \longrightarrow P_2 \otimes_R K \xrightarrow{d_2 \otimes 1} P_1 \otimes_R K \xrightarrow{d_1 \otimes 1} P_0 \otimes_R K \longrightarrow 0$$

Then, we can take the homology groups of our new complex $P_\bullet \otimes_R K$; these are our Tor groups.

$$\mathrm{Tor}_n^R(N, K) := H_n(P_\bullet \otimes_R K) = \ker(d_n \otimes 1)/\mathrm{im}(d_{n+1} \otimes 1)$$

Moreover, these Tor groups are independent of the choice of projective resolution.

Again, the proof that Tor_n^R is independent of the choice of projective resolution requires more machinery than we have developed.

Remark 6.8. Recall that since we are working over commutative rings R, $M \otimes_R N = N \otimes_R M$. It follows that $\mathrm{Tor}_n^R(M, N) \cong \mathrm{Tor}_n^R(N, M)$.

Now, we can have an analogue of Proposition 5.20 for $\mathrm{Tor}_n^R(N, M)$, involving flat modules.[6]

Proposition 6.9. *Let M be an R-module. The following statements are equivalent. Moreover, due to the symmetry of* Tor*, we may make the same statements for N.*

(1) *M is flat.*
(2) $\mathrm{Tor}_1^R(N, M) = 0$ *for all R-modules N.*
(3) $\mathrm{Tor}_n^R(N, M) = 0$ *for all R-modules N and $n > 0$.*

So, we can see that Tor measures the failure of a module to be flat, just like how $\mathrm{Ext}_n^R(M, N)$ measure the failure of M to be projective (or N injective). At this point, you must be wondering if there's a nice and clean intuitive view of Tor_1, just like with Ext_1. Well, you'd be in luck!

Definition 6.10. Let R be an integral domain, and let M be an R-module. The *torsion submodule*[7] of M, denoted M_T, is the following R-submodule.

$$M_T = \{m \in M \mid \exists r \in R \setminus \{0\} \text{ such that } rm = 0\}$$

The module M is said to be *torsion-free* if $M_T = 0$, or M is a *torsion* module if $M_T = M$.

This is a generalization of the torsion subgroup of an abelian group, which is the subgroup of elements of finite order.[8]

The following theorem will give us some intuition for Tor_1.

Theorem 6.11. *Let R be an integral domain with field of fractions K, and let M be an R-module. Then,* $\mathrm{Tor}_1^R(K/R, M)$ *is isomorphic to M_T.*

Proposition 6.12. *A flat module over an integral domain is torsion-free.*

Proof. Let M be a flat R-module, with R an integral domain. Then, $\mathrm{Tor}_1^R(N, M) = 0$ for all R-modules N. Since K/R is an R-module, with K the field of fractions of R, it follows that $\mathrm{Tor}_1^R(K/R, M) = 0$. But, by the previous theorem, $M_T = 0$ and thus M is torsion-free. $\square$

[5] For no particular reason, of course.
[6] We now return to normal notation.
[7] This is where "Tor" comes from.
[8] In fact, this definition coincides with that from abelian groups when we consider $R = \mathbb{Z}$.

References

[Cla] Pete L Clark. *Math 8030 Introduction To Homological Algebra*. URL: http://math.uga.edu/~pete/Math8030_Exercises.pdf.

[CS] Wojciech Chacholski and Roy Skjelnes. *Homological Algebra and Algebraic Topology*. URL: https://www.math.kth.se/math/GRU/2011.2012/SF2735/Notesforsf2735.pdf.

[Gun] Kate Gunzinger. *It's My Turn (1980) Snake Lemma*. URL: https://www.youtube.com/watch?v=etbcKWEKnvg.

[Kob] Andrew Kobin. *Homological Algebra*. URL: https://static1.squarespace.com/static/5aff705c5ffd207cc87a512d/t/5c096a34352f53aaa565c6d5/1544120889673/Homological+Algebra.pdf.

[Lin] Markus Linckelmann. *Basic Homological Algebra*. URL: http://www.ltcc.ac.uk/media/london-taught-course-centre/documents/Background-Notes.pdf.

[May] Peter May. *Notes on Tor and Ext*. URL: http://math.uchicago.edu/~may/MISC/TorExt.pdf.

[Rot09] Joseph J. Rotman. *An Introduction to Homological Algebra*. New York, NY: Springer New York, 2009. ISBN: 978-0-387-24527-0 978-0-387-68324-9. DOI: 10.1007/b98977. URL: http://link.springer.com/10.1007/b98977.

[Sat] Sean Sather-Wagstaff. *Homological Algebra Notes*. URL: https://www.ndsu.edu/pubweb/~ssatherw/DOCS/han.pdf.