

Mathematics 6310
Introduction to category theory
Ken Brown, Cornell University, September 2010

1. INTRODUCTION

Mathematical objects don't exist in isolation. To understand a group G , for example, one must often consider homomorphisms between G and other groups. The class of groups, together with all homomorphisms between them, forms the "category" of groups. Here are a few other familiar categories:

- The category of topological spaces and continuous maps.
- The category of vector spaces and linear maps (over a given field).
- The category of sets and (arbitrary) functions.
- The category of posets and order-preserving functions.

Although category theory can get quite abstract, it leads to a way of thinking that I find very useful. In this handout I will introduce a few notions of category theory that will occur over and over again in the course. You can supplement this handout by reading Appendix II in your text, starting on p. 911, but the handout is intended to be self-contained.

2. CATEGORIES

A *category* \mathcal{C} consists of the following:

- a class $\text{ob } \mathcal{C}$ of *objects*;
- for each pair of objects A, B , a set $\text{Hom}(A, B)$, or $\text{Hom}_{\mathcal{C}}(A, B)$, of *morphisms*; and
- for each triple of objects A, B, C , a *composition* operation

$$\text{Hom}(A, B) \times \text{Hom}(B, C) \rightarrow \text{Hom}(A, C),$$

denoted $(f, g) \mapsto gf$ or $g \circ f$.

We require composition to be associative in the obvious sense: $(hg)f = h(gf)$ whenever the composites are defined. We also require the existence of identities id_A , one for each object A , such that

$$f \circ \text{id}_A = f = \text{id}_B \circ f$$

for any $f \in \text{Hom}(A, B)$.

It is clear what the objects and morphisms are in the categories described informally in Section 1. But the abstract definition also leads to some new examples. For example, a category with one object is essentially the same as a monoid (i.e., semigroup with identity). Posets also provide examples. If X is a poset, then there is a category with X as its set of objects and with a unique morphism $x \rightarrow y$ whenever $x \leq y$ in X . These examples are pretty concrete. A more abstract example (the dual of a category) will be given in the next section.

It is often easier to understand arguments in category theory if we use diagrams rather than equations. Thus we will often write

$$f: A \rightarrow B$$

to mean $f \in \text{Hom}(A, B)$. The notion of “commutative diagram” then makes sense in an arbitrary category. For example, an equation $h = gf$ can be expressed by a commutative triangle of the form

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ & \searrow h & \downarrow g \\ & & C \end{array}$$

Note that the concept of isomorphism makes sense in every category: A morphism $f: A \rightarrow B$ is an *isomorphism* if it has an inverse $g: B \rightarrow A$ (i.e., a morphism such that $fg = \text{id}_B$ and $gf = \text{id}_A$). The inverse is then unique by the standard argument that you’ve seen in other settings (which only requires associativity of composition).

3. DUALITY

Associated to every category \mathcal{C} is its *dual* category \mathcal{C}^{op} , with the same objects but with “all arrows reversed”. Thus

$$\text{Hom}_{\mathcal{C}^{\text{op}}}(A, B) := \text{Hom}_{\mathcal{C}}(B, A).$$

Composition in \mathcal{C}^{op} is defined to be the opposite of composition in \mathcal{C} . In other words, if gf is defined in \mathcal{C} , we write fg for the same morphism viewed as a morphism in \mathcal{C}^{op} . In terms of diagrams, this simply means that a commutative diagram in \mathcal{C}^{op} is exactly the same thing as a commutative diagram in \mathcal{C} , but with all arrows reversed.

If this seems confusing, you might find the following notational convention useful: Given an object A of \mathcal{C} , write A^* for the same object viewed as an object of \mathcal{C}^{op} . And given a morphism $f: A \rightarrow B$ in \mathcal{C} , write $f^*: B^* \rightarrow A^*$ for the same morphism viewed as a morphism in \mathcal{C}^{op} . The definition of composition in \mathcal{C}^{op} then becomes

$$f^*g^* := (gf)^*$$

whenever the composites are defined. The notation is intended to remind you of duality in vector space theory.

This abstract construction is extremely useful. It means that every concept about categories has a dual concept, with all arrows reversed. Formally, we just apply the original concept to the dual category. Similar remarks apply to theorems about general categories.

4. FUNCTORS

Categories don’t exist in isolation. To understand a category \mathcal{C} , one must often consider “functors” between \mathcal{C} and other categories. A functor is something like a “morphism between categories”. Algebraic topology, for example, provides many functors from topological spaces to groups (homology, fundamental group, . . .). Indeed, category theory grew out of the study of such functors.

A *functor* $T: \mathcal{C} \rightarrow \mathcal{D}$ consists of two functions. First, T associates to each object A of \mathcal{C} an object $T(A)$ of \mathcal{D} . Second, T associates to each morphism $f: A \rightarrow B$ in \mathcal{C} a morphism $T(f): T(A) \rightarrow T(B)$ in \mathcal{D} . This function on morphisms is required

to preserve identities and compositions (like a homomorphism of monoids). Thus $T(\text{id}_A) = \text{id}_{T(A)}$, and $T(fg) = T(f)T(g)$ whenever the composition fg is defined.

When describing a functor, we often just say what it does to objects, since the effect on morphisms is often obvious. For example, there is a functor that associates to every group G its abelianization $G/[G, G]$; it should be easy for you to figure out how a group homomorphism induces a homomorphism on abelianizations.

Functors, as just defined, are more properly called *covariant* functors, since the arrows in \mathcal{D} go in the same direction as the arrows in \mathcal{C} . There is also a notion of *contravariant* functor, in which a morphism $f: A \rightarrow B$ gives rise to a morphism $T(f): T(B) \rightarrow T(A)$ in the opposite direction. The effect on composition, then, is that $T(fg) = T(g)T(f)$, i.e., T reverses compositions instead of preserving them.

Note that a contravariant functor $\mathcal{C} \rightarrow \mathcal{D}$ is the same thing as a covariant functor $\mathcal{C}^{\text{op}} \rightarrow \mathcal{D}$, so there is no need to develop a separate theory of contravariant functors. In other words, every result about covariant functors can be dualized to give a result about contravariant functors.

The most familiar example of a contravariant functor is the duality functor on vector spaces, which associates to every vector space V its dual space $V^* := \text{Hom}(V, k)$, where k is the field of scalars. Note that a linear transformation $f: V \rightarrow W$ induces a homomorphism $f^*: W^* \rightarrow V^*$, so this functor is contravariant. If we choose bases and represent linear transformations by matrices, the matrix of f^* is the transpose of the matrix of f . The law $(fg)^* = g^*f^*$ then reduces to the familiar fact that transposition of matrices reverses products.

Functors, in what follows, will always be assumed covariant unless the contrary is explicitly stated or is clear from the context.

5. NATURALITY

One of the goals of the developers of category theory in the 1940s was to give a precise meaning to the word “natural”, which is often used informally. The definition they came up with has proven extremely useful, although it takes some experience before one can appreciate that it (usually) captures the informal meaning of “natural”.

Suppose we are given two functors $S, T: \mathcal{C} \rightarrow \mathcal{D}$ and, for each object A of \mathcal{C} , a morphism $\phi(A): S(A) \rightarrow T(A)$ in \mathcal{D} . We say that this collection ϕ of morphisms is *natural*, or is a *natural transformation* from S to T , if it satisfies the following condition: For every morphism $f: A \rightarrow B$ in \mathcal{C} , the square

$$(1) \quad \begin{array}{ccc} S(A) & \xrightarrow{S(f)} & S(B) \\ \phi(A) \downarrow & & \downarrow \phi(B) \\ T(A) & \xrightarrow{T(f)} & T(B) \end{array}$$

commutes (in \mathcal{D}). A natural transformation ϕ is said to be a *natural isomorphism* if it has an inverse natural transformation ψ . This is the same as saying that $\phi(A)$ is an isomorphism in \mathcal{D} for each A [check this].

A familiar example from linear algebra is that a finite-dimensional vector space V is naturally isomorphic to its double dual V^{**} . The precise statement is that the identity functor on the category of finite-dimensional vector spaces is naturally

isomorphic to the functor $V \mapsto V^{**}$. Note that this makes sense, since the second functor is covariant, being the composite of two contravariant functors. But it does *not* make sense to ask whether the identity functor is naturally isomorphic to the contravariant functor $V \mapsto V^*$.


For practice, you should check that the canonical map $V \rightarrow V^{**}$ really gives a natural transformation of functors in the technical sense defined above. It's easy, but you do have to work through the square (1) and remember the definitions.

6. SUMS AND PRODUCTS

Let X_1 and X_2 be sets, and consider their Cartesian product $X = X_1 \times X_2$. If Y is any other set, then a function $f: Y \rightarrow X$ is determined by its two coordinate functions $f_1: Y \rightarrow X_1$ and $f_2: Y \rightarrow X_2$. Thus we have a bijection of sets

$$(2) \quad \text{Hom}(Y, X) \xrightarrow{\sim} \text{Hom}(Y, X_1) \times \text{Hom}(Y, X_2)$$

given by $f \mapsto (f_1, f_2)$. Note that we can recover the coordinate functions f_i from f by $f_i = f\pi_i$, where $\pi_i: X \rightarrow X_i$ is the canonical projection ($i = 1, 2$). We can summarize the notation in the following diagram:



Suppose now that the sets have some algebraic structure. For example, they may be groups, rings, vector spaces, In many familiar cases we can impose the same structure on $X_1 \times X_2$ by defining the operations componentwise, and everything above remains valid in the corresponding category. For example, if X_1 and X_2 are groups, then X is a group, and equation (2) is still valid if “Hom” refers to group homomorphisms. Another example that you may have seen is the product of two topological spaces. Here the set-theoretic product is given a topology, which is cooked up precisely so that a map into the product is continuous if and only if its component functions are continuous.

From the point of view of category theory, we have captured the essence of what it means to be a product in each of the examples above by spelling out what it means to give a morphism into a product. We now generalize to an arbitrary category.

Let X_1 and X_2 be objects in a category \mathcal{C} . Suppose there is an object X that admits morphisms $\pi_1: X \rightarrow X_1$ and $\pi_2: X \rightarrow X_2$ with the following universal property: Given any object Y and any morphisms $f_1: Y \rightarrow X_1$ and $f_2: Y \rightarrow X_2$, there is a unique morphism $f: Y \rightarrow X$ such that $\pi_1 f = f_1$ and $\pi_2 f = f_2$, as in the diagram above. A simple argument shows that X is then unique up to canonical isomorphism. [You may have seen such arguments in other settings. If not, and if you can't figure it out, please ask.] We write $X = X_1 \times X_2$ and call it the *product* of X_1 and X_2 . The UMP characterizing the product can then be summarized as a

bijection of sets

$$\mathrm{Hom}_{\mathcal{C}}(Y, X_1 \times X_2) \xrightarrow{\sim} \mathrm{Hom}_{\mathcal{C}}(Y, X_1) \times \mathrm{Hom}_{\mathcal{C}}(Y, X_2)$$

given by composition with π_1 and π_2 .

All of this extends in an obvious way to the product $\prod_{j \in J} X_j$ of an arbitrary family of objects. We dealt with two objects above just to keep the notation simple.

By reversing all the arrows (or working in the dual category), we can define the *sum* (also called the *coproduct*) $X = \coprod_{j \in J} X_j$ of a family of objects X_j . It comes equipped with morphisms $i_j: X_j \rightarrow X$. The UMP (in the case $X = X_1 \amalg X_2$ for simplicity) takes the form

$$\begin{array}{ccc}
 X_1 & \xrightarrow{f_1} & Y \\
 \searrow i_1 & & \uparrow f \\
 & X \dashrightarrow & \\
 \nearrow i_2 & & \uparrow f_2 \\
 X_2 & \xrightarrow{f_2} & Y
 \end{array}$$

And we have a bijection of sets

$$(3) \quad \mathrm{Hom}_{\mathcal{C}}(X_1 \amalg X_2, Y) \xrightarrow{\sim} \mathrm{Hom}_{\mathcal{C}}(X_1, Y) \times \mathrm{Hom}_{\mathcal{C}}(X_2, Y).$$

given by composition with i_1 and i_2 .

Sums occur naturally when one wants to “glue” two objects together. In the category of sets, the sum is the disjoint union, with the i_j being inclusion maps. In the category of topological spaces, the sum is the set-theoretic disjoint union, topologized so that each summand is an open subspace. In the category of vector spaces (or abelian groups), the sum is the direct sum $\bigoplus_{j \in J} X_j$. The sum in the category of groups is a less familiar object, called the *free product*, which we will study later. (This confusing terminology predates category theory and is not likely to change.) There is a similar construction for rings. The sum in the category of commutative rings requires tensor products; again, we will study it later.

7. REPRESENTABLE FUNCTORS AND YONEDA’S LEMMA

Let \mathcal{C} be an arbitrary category. The very definition of “category” gives us a whole slew of functors $h_A := \mathrm{Hom}(A, -)$ from \mathcal{C} to the category of sets, one for each object A of \mathcal{C} . The functor

$$h_A: \mathcal{C} \rightarrow (\text{sets})$$

is given on objects by

$$h_A(B) := \mathrm{Hom}_{\mathcal{C}}(A, B)$$

and on morphisms by composition. More precisely, given $f: B \rightarrow C$,

$$h_A(f): h_A(B) \rightarrow h_A(C)$$

sends $g \in h_A(B) = \mathrm{Hom}(A, B)$ to $fg \in h_A(C) = \mathrm{Hom}(A, C)$. [You should check that this definition really makes h_A a functor.] These functors h_A are called *representable* functors.

Dually, each object A of \mathcal{C} gives rise to a contravariant representable functor

$$h^A: \mathcal{C} \rightarrow (\text{sets})$$

given by

$$h^A := \text{Hom}_{\mathcal{C}}(-, A).$$

Many of the UMPs that you have seen can be viewed as descriptions of certain representable functors. For example, the UMP for a quotient group G/N describes the representable functor $h_{G/N}$. [It is equivalent to a certain subfunctor of h_G .] And the UMP for a sum $X := X_1 \amalg X_2$ describes the representable functor h_X . [It is equivalent to $h_{X_1} \times h_{X_2}$.] Another example that you saw in a homework problem involved the representable functor $h_{G/H}$ in the category of G -sets. [It is equivalent to the functor $X \mapsto X^H$.]

Sometimes it's possible to endow the sets $h_A(B)$ or $h^A(B)$ with additional structure, so that the representable functors can be viewed as taking values in some subcategory of the category of sets. For example, the functor $V \mapsto V^*$ on k -vector spaces is a contravariant representable functor in which the sets $V^* = \text{Hom}(V, k)$ have been given the structure of k -vector space. Similarly, all the representable functors on the category of R -modules (for an arbitrary commutative ring R) can be viewed as taking values in the category of R -modules. If R is noncommutative, they can be viewed as taking values in the category of abelian groups.

The definition of “category” not only gives us the functors h_A , but it also gives us canonical elements $\text{id}_A \in h_A(A)$. The following result, called *Yoneda's lemma*, can be viewed as saying that the functor h_A is “freely generated” by the element id_A .

Proposition 1. *Let $T: \mathcal{C} \rightarrow (\text{sets})$ be an arbitrary functor, and let A be an object of \mathcal{C} . For any element $u \in T(A)$, there is a unique natural transformation*

$$\phi: h_A \rightarrow T$$

such that $\phi(A): h_A(A) \rightarrow T(A)$ takes id_A to u . For any object B of \mathcal{C} , the map $\phi(B): h_A(B) \rightarrow T(B)$ is given by

$$(4) \quad \phi(B)(f) = T(f)(u)$$

for $f \in h_A(B) = \text{Hom}(A, B)$.

Proof. The formula (4) (and hence the uniqueness of ϕ) follows immediately from the diagram

$$\begin{array}{ccc} h_A(A) & \xrightarrow{h_A(f)} & h_A(B) \\ \phi(A) \downarrow & & \downarrow \phi(B) \\ T(A) & \xrightarrow{T(f)} & T(B) \end{array}$$

defining “natural transformation”. For existence, take (4) as a definition, and check naturality. \square

As a special case, suppose that T is also a representable functor h_B . Yoneda's lemma then gives:

Corollary. *Natural transformations $h_A \rightarrow h_B$ are in 1-1 correspondence with morphisms $B \rightarrow A$. In particular, h_A is naturally equivalent to h_B if and only if $A \cong B$. \square*

The second assertion formalizes the informal statement that if an object is characterized by a UMP, then it is unique up to isomorphism.

As another illustration, suppose $X = X_1 \amalg X_2$. Then each component of the bijection in (3) is a map of the type defined in (4) (and even of the type that occurs in the corollary), so we conclude that (3) is natural when both sides are viewed as functors of Y . Conversely, suppose we are given objects X_1, X_2, X and a natural isomorphism

$$\mathrm{Hom}_{\mathcal{C}}(X, Y) \xrightarrow{\cong} \mathrm{Hom}_{\mathcal{C}}(X_1, Y) \times \mathrm{Hom}_{\mathcal{C}}(X_2, Y).$$

Then Yoneda's lemma (or even the corollary) implies that the isomorphism is induced by maps $i_1: X_1 \rightarrow X$ and $i_2: X_2 \rightarrow X$, so that $X = X_1 \amalg X_2$. Similar remarks apply to arbitrary sums and, by duality, to products.

8. ADJOINT FUNCTORS

Let \mathcal{C} and \mathcal{D} be categories, and suppose we are given functors

$$\mathcal{C} \begin{array}{c} \xrightarrow{S} \\ \xleftarrow{T} \end{array} \mathcal{D}$$

We say that S is a *left adjoint* of T and that T is a *right adjoint* of S if there is a natural isomorphism

$$(5) \quad \mathrm{Hom}_{\mathcal{D}}(S(C), D) \cong \mathrm{Hom}_{\mathcal{C}}(C, T(D))$$

for $C \in \mathrm{ob} \mathcal{C}$ and $D \in \mathrm{ob} \mathcal{D}$. Here both sides of (5) are viewed as set-valued functors of the two variables C, D , contravariant in C and covariant in D . Adjoint functors are ubiquitous; once you are aware of them, you will notice them everywhere.

Here is a simple example. For any set A , let $F(A)$ be the free group generated by A . Then the UMP for free groups gives us a bijection

$$\mathrm{Hom}_{\mathrm{groups}}(F(A), G) \cong \mathrm{Hom}_{\mathrm{sets}}(A, G)$$

for any group G . This is easily checked to be natural, and it shows that the functor $F: (\mathrm{sets}) \rightarrow (\mathrm{groups})$ is left adjoint to the forgetful functor $(\mathrm{groups}) \rightarrow (\mathrm{sets})$. (The forgetful functor associates to a group G the underlying set G , i.e., it “forgets” that G has a group structure.)

A second example from group theory involves the *abelianization functor* $G \mapsto G_{\mathrm{ab}}$, where the latter is the quotient $G/[G, G]$. If we view this as a functor from groups to abelian groups, then it is left adjoint to the forgetful functor from abelian groups to groups. (In this case the forgetful functor takes an abelian group to itself, just viewed as a group; in other words, it “forgets” that the group is abelian.)

Suppose we are given an adjunction

$$(6) \quad \phi = \phi(C, D): \mathrm{Hom}_{\mathcal{D}}(S(C), D) \xrightarrow{\cong} \mathrm{Hom}_{\mathcal{C}}(C, T(D))$$

as in (5). If we fix C and view both sides as functors of D , then Yoneda's lemma shows that $\phi(C, -)$ is induced by an element of $\mathrm{Hom}_{\mathcal{C}}(C, T(S(C)))$, i.e., by a morphism

$$\eta(C): C \rightarrow T(S(C))$$

in \mathcal{C} . We obtain $\eta(C)$ from ϕ by putting $D = S(C)$ in (6) and taking the image of $\text{id}_{S(C)}$. And we recover ϕ from the maps $\eta(C)$ as follows: Given f in $\text{Hom}_{\mathcal{D}}(S(C), D)$, its image $\phi(f)$ in $\text{Hom}_{\mathcal{C}}(C, T(D))$ is the composite

$$C \xrightarrow{\eta(C)} T(S(C)) \xrightarrow{T(f)} T(D).$$

In our free group example, $\eta(A)$ is simply the canonical set-theoretic inclusion $A \hookrightarrow F(A)$. In the abelianization example, $\eta(G)$ is the quotient map $G \twoheadrightarrow G_{\text{ab}}$.

Similarly, the inverse ψ of ϕ is necessarily induced by a family of morphisms

$$\varepsilon(D): S(T(D)) \rightarrow D,$$

one for each $D \in \text{ob } \mathcal{D}$. In our free group example, this is the canonical surjection $F(G) \twoheadrightarrow G$ that we would use if we wanted to prove, without an arbitrary choice of generators, that every group G is a quotient of a free group. In the abelianization example, $ST(A) = A/\{1\}$ for an abelian group A , and $\varepsilon(A)$ is the natural isomorphism $A/\{1\} \rightarrow A$. [Here the word “natural” can be interpreted either formally or informally.]

The fact that ϕ and ψ are actually inverses of one another is reflected in the fact that we get identity morphisms by forming the “obvious” composites

$$S(C) \rightarrow S(T(S(C))) \rightarrow S(C)$$

and

$$T(D) \rightarrow T(S(T(D))) \rightarrow T(D).$$

I leave it as an exercise for you to figure out where all these arrows come from.

Note that we haven’t yet made use of the fact that the adjunction isomorphisms are supposed to be natural when both sides of (6) are viewed as functors of *two* variables. The translation of this is that the families $\eta := (\eta(C))$ and $\varepsilon := (\varepsilon(D))$ are required to be natural transformations (in \mathcal{C} and \mathcal{D} , respectively).

To summarize what we have done so far, an adjunction between S and T is determined by natural transformations $\eta: \text{id}_{\mathcal{C}} \rightarrow TS$ and $\varepsilon: ST \rightarrow \text{id}_{\mathcal{D}}$ satisfying certain identities (which say that the composites above are identity morphisms).

Our next observation is that each of the functors S, T determines the other one up to natural isomorphism. So if a functor has a left or right adjoint, the latter is essentially unique. Here is a sketch of the proof. Suppose $T: \mathcal{D} \rightarrow \mathcal{C}$ has two left adjoints S, S' . Then we have a natural isomorphism

$$(7) \quad \text{Hom}_{\mathcal{D}}(S(C), D) \cong \text{Hom}_{\mathcal{D}}(S'(C), D).$$

For each fixed C , this is an isomorphism between two representable functors on \mathcal{D} , so it is induced by a unique isomorphism $S(C) \cong S'(C)$. This is natural in C because (7) is natural in C , so S is naturally isomorphic to S' .

Finally, we remark that if S and T are adjoints as above, then S preserves sums and (dually) T preserves products. Here’s a sketch of the proof. If $X = \coprod_{j \in J} X_j$, then we have isomorphisms (natural in Y)

$$\begin{aligned} \text{Hom}(S(X), Y) &\cong \text{Hom}(X, T(Y)) \\ &\cong \prod_{j \in J} \text{Hom}(X_j, T(Y)) \\ &\cong \prod_{j \in J} \text{Hom}(S(X_j), (Y)). \end{aligned}$$

Now apply the observations at the end of Section 7 to deduce that

$$S(X) = \prod_{j \in J} S(X_j).$$

9. GALOIS CONNECTIONS

If we specialize to categories that arise from posets, a pair of adjoint functors is called a *Galois connection*. Let X and Y be posets, viewed as categories as explained in Section 2. Since there is at most one morphism between two given objects, several concepts of category theory simplify drastically:

- A functor is simply an order-preserving map.
- If there is a natural transformation from one such map to another, it is unique.
- Naturally isomorphic maps are equal.

Keeping these simplifications in mind, we see that a pair of adjoint functors between X and Y consists of a pair of order-preserving maps

$$X \begin{array}{c} \xrightarrow{S} \\ \xleftarrow{T} \end{array} Y$$

such that

$$S(x) \leq y \iff x \leq T(y)$$

for $x \in X$ and $y \in Y$.

Exercise 1. Recast the definition in terms of the ε, η characterization of adjoint functors.

The remarks at the end of the previous section show that each of the maps S, T determines the other. Moreover, S preserves least upper bounds (categorical sums) and T preserves greatest lower bounds (categorical products).

Before giving examples, we switch to the case where the maps S, T are order-reversing instead of order-preserving, since this is the case in classical examples. We can reduce this case to the previous one by replacing Y by Y^{op} . The adjunction condition then becomes

$$y \leq S(x) \text{ in } Y \iff x \leq T(y) \text{ in } X.$$

Note that there is now complete symmetry between S and T , so there is no need to distinguish between “left” and “right”.

Exercise 2. Rewrite the result of Exercise 1 for order-reversing maps.

From now on all Galois connections will be assumed to be order reversing.

Example 1. Let A and B be sets, and let R be a relation between them. Then there is a Galois connection between subsets of A (ordered by inclusion) and subsets of B (ordered by inclusion), given by

$$S(E) := \{b \in B \mid aRb \text{ for all } a \in E\}$$

for $E \subseteq A$, and

$$T(F) := \{a \in A \mid aRb \text{ for all } b \in F\}$$

for $F \subseteq B$. Note that we have

$$F \subseteq S(E) \iff aRb \text{ for all } a \in E \text{ and } b \in F \iff E \subseteq T(F),$$

so S and T define a Galois connection.

Example 2. Suppose a group G acts on a set A . Define a relation between G and A by

$$gRa : \iff ga = a$$

for $g \in G$ and $a \in A$. Then the previous example yields a Galois connection between subsets of G and subsets of A given by

$$S(E) := A^E$$

for $E \subseteq G$, and

$$T(F) := \text{Fix}(F)$$

for $F \subseteq A$, where $\text{Fix}(F)$, the “fixer” of F , is the set of group elements that fix F pointwise. (It is generally a proper subgroup of the stabilizer of F , which consists of the group elements that fix F as a set.) This is the classical situation of Galois theory, in which A is a field and G is a finite group of automorphisms of A . It is this example that is responsible for the name “Galois connection”.

Example 3. Let k be a field, let $A := k^n$ (“affine space”), and let R be the ring of polynomial functions on A . It can be identified with the polynomial ring $k[x_1, \dots, x_n]$ if k is infinite; otherwise it is a quotient of $k[x_1, \dots, x_n]$. Define a relation between R and A by

$$fRa : \iff f(a) = 0$$

for $f \in R$ and $a \in A$. Then Example 1 yields a Galois connection between subsets of R and subsets of A , given by

$$S(E) := \mathcal{Z}(E) := \{a \in A \mid f(a) = 0 \text{ for all } f \in E\}$$

for $E \subseteq R$, and

$$T(F) := \mathcal{I}(F) := \{f \in R \mid f(a) = 0 \text{ for all } a \in F\}$$

for $F \subseteq A$. We call $\mathcal{Z}(E)$ the zero set of E and $\mathcal{I}(F)$ the ideal vanishing on F . We will return to this example in the algebraic geometry portion of the course.

Remark. Given a Galois connection between two posets X, Y , it follows from Exercise 2 that S and T define mutually inverse bijections between the image of T and the image of S . To give this some content, one needs to determine these images. Consider the case of classical Galois theory (Example 2, with G a finite group of automorphisms of a field K). Then the image of S is the set of all fields intermediate between K and the fixed field K^G , while the image of T is the set of all subgroups of G . The bijection between these sets is called the fundamental theorem of Galois theory. Or consider the case of classical algebraic geometry (Example 3, with k algebraically closed). Then the image of S is the set of all algebraic subvarieties of A [by definition], while the image of T is the set of radical ideals in R . This last statement is one version of a famous theorem called Hilbert’s Nullstellensatz, which we will prove.