Lecture 1: Overview

January 24, 2018

We begin with a very quick review of first-order logic (we will give a more leisurely review in the next lecture).

Recall that a *linearly ordered set* is a set X equipped with a binary relation $R(x, y)$ (usually written as $x \leq y$) which satisfies the following axioms:

$(A1)$ $(\forall x)[R(x, x)]$

$(A2)$ $(\forall x, y, z)[(R(x, y) \text{ and } R(y, z)) \Rightarrow R(x, z)]$

$(A3)$ $(\forall x, y)[(R(x, y) \text{ and } R(y, x)) \Rightarrow (x = x)]$

$(A4)$ $(\forall x, y)[R(x, y) \text{ or } R(y, x)]$

Axioms $(A1)$ through $(A4)$ are examples of *first-order sentences*: they are statements that can be built in finitely many steps from the relation symbol R and the binary relation $=$ using basic logical operations ("and", "or", "implies", "not") and quantification ("for all", "there exists") over the domain X.

Definition 1. A *language* L consists of a set of *predicate symbols* $\{P_i\}_{i \in I}$, each of which has some specified *arity* $n_i \geq 0$.

A *first-order theory* T consists of a language L together with a collection of first-order sentences which are constructed using the symbols of L; we will refer to those sentences as the *axioms of* T.

Example 2. Let L be the language consisting of a single predicate symbol R of arity 2. Then L, together with the collection of first-order sentences $(A1)$ through $(A4)$ above, comprise a first-order theory T. We will refer to T as the *theory of linear orderings*.

Definition 3. Let L be a language consisting of predicate symbols $\{P_i\}_{i \in I}$ of arity n_i. An *L-structure* is a set M together with a subset $M[P_i] \subseteq M^{n_i}$ for each predicate symbol P_i; we think of $M[P_i]$ as the set of all tuples $(c_1, c_2, \ldots, c_{n_i}) \in M^{n_i}$ which satisfy the predicate P_i.

Given an L-structure M and a first-order sentence φ in the language L, we can ask if φ is *true in* M; in this case, we write $M \vDash \varphi$.

Let T be a first-order theory with underlying language L. A *model of* T is an L-structure M such that $M \vDash \varphi$ for each axiom φ of T.

Example 4. Let T be the theory of linear orderings. Then a model of T is a set M together with a binary relation R satisfying axioms $(A1)$ through $(A4)$: that is, a set with a linear ordering.

Example 5 (Abelian Groups)**.** An abelian group can be understood as a set A with a ternary relation P, where $P(x, y, z) \Leftrightarrow (x + y = z)$. This relation must be required to satisfy the following (first-order) axioms:

$(B1)$

$$(\forall x, y)(\exists z)[P(x, y, z)]$$

$(B2)$

$$(\forall x, y, z', z')[(P(x, y, z) \text{ and } P(x, y, z')) \Rightarrow (z = z')]$$

($B3$)
$$(\forall x, y, z)[P(x, y, z) \Rightarrow P(y, x, z)]$$

($B4$)
$$(\exists e)[((\forall x)P(x, e, x)) \text{ and } (\forall x)(\exists y)P(x, y, e)]$$

($B5$)
$$(\forall x, y, z, u, v, w)[P(x, y, u) \text{ and } P(u, z, w) \text{ and } P(y, z, v) \Rightarrow P(x, v, w)].$$

Here ($B1$) and ($B2$) assert that $P(x, y, z)$ determines an addition $+ : A \times A \to A$, ($B3$) guarantees that addition is commutative, ($B4$) guarantees that there is an identity element and additive inverses, and ($B5$) encodes the associative law.

We can therefore view abelian groups as the models of a first-order theory T, whose underlying language has a single 3-ary relation symbol P, and whose axioms are ($B1$) through ($B5$). We refer to T as *the theory of abelian groups.*

One can give endless variants of Example 5: groups, rings, fields, ordered fields, and so forth can all be understood as models of a suitably chosen first-order theory. However, not all mathematical structures of interest can be described this way. For example, there is no first order theory of torsion groups, or of finite fields, or of well-ordered sets. The expressive power of first-order logic is deliberately limited, in part because we want to be able to prove results like this one:

Theorem 6 (Gödel's Completeness Theorem)**.** *Let T be a first-order theory and let φ be a sentence in the language of T. The following conditions are equivalent:*

(1) *There is a proof of φ from the axioms of T (written as $T \vdash \varphi$).*

(2) *The sentence φ is true in every model of T (written as $T \vDash \varphi$).*

Of course, to make Theorem 6 precise, we need to define what it means to give a proof of φ from the axioms of T. We will return to this point later. For the moment, let us just note that any reasonable notion of proof should be something that is *algorithmically verifiable.* So, for example, if T is the theory of abelian groups (Example 5), then it should be possible to program a computer to generate the (infinite) list of all first-order sentences (in the language $L = \{P\}$) which are provable from T. By virtue of Theorem 6, this is the same as the list of all first-order sentences (in the language L) which are true in all abelian groups. It is not at all obvious that this latter list should be computably enumerable (even very fast computers will have trouble checking something in every abelian group).

Warning 7. In the setting of Theorem 6, it is essential that we restrict our attention to *first-order* sentences. There is no counterpart of Theorem 6 for more expressive languages, like second-order logic. For example, we can characterize $\mathbf{Z}$ as the unique linearly ordered ring in which the set of positive elements $\{x : x \geq 0\}$ is well-ordered. The axiom that $\{x : x \geq 0\}$ is well-ordered is a *second-order* statement: it involves quantification over subsets of $\mathbf{Z}$, rather than over elements of $\mathbf{Z}$. If there was a reasonable analogue of Theorem 6 for second-order theories, then it would be possible to program a computer to generate a list of all statements which are true about $\mathbf{Z}$ (as a linearly ordered ring). However, it is known that such a computer program cannot exist (even if we restrict our attention to first order sentences).

Theorem 6 establishes a relationship between *syntax* and *semantics.* The relation $T \vdash \varphi$ is syntactic, in the sense that it is defined in terms of formal manipulation of symbols. On the other hand, the relation $T \vDash \varphi$ is semantic: it involves thinking about the models of T, and how φ is interpreted in those models. To bring the contrast into sharper relief, it will be convenient to state an equivalent form of Theorem 6. First, we need a bit of terminology: given a language L, a first-order *formula* is a statement $\varphi(x_1, \ldots, x_n)$ that can be formulated in the language L, but makes reference to certain "free" variables. For example, if $L = \{R\}$ is the language of linear orderings, then $(\forall y)[R(x, y)]$ is a formula with a single free variable x, which asserts that x is the least element of the linear order in question.

Given a formula $\varphi(x_1, \ldots, x_n)$ with n free variables and a model M of T, we can define a set

$$M[\varphi] = \{(c_1, c_2, \ldots, c_n) \in M^n : \ \varphi(c_1, \ldots, c_n) \text{ is true in } M \ \} \subseteq M^n.$$

Theorem 8 (Gödel's Completeness Theorem, Equivalent Form)**.** *Let T be a first-order theory and suppose we are given a pair of formulas $\varphi(x_1, \ldots, x_n)$, $\varphi'(x_1, \ldots, x_n)$ in the language of T. The following conditions are equivalent:*

(1) *There is a proof, from the axioms of T, that $\varphi(x_1, \ldots, x_n)$ and $\varphi'(x_1, \ldots, x_n)$ are equivalent.*

(2) *For every model M of T, the sets $M[\varphi]$ and $M[\varphi']$ coincide (as subsets of M^n).*

Once again, Theorem 8 is about the relationship between syntax and semantics. We can define an equivalence relation on the set of formulas in some fixed number of variables, where we declare $\varphi(x_1, \ldots, x_n)$ and $\varphi'(x_1, \ldots, x_n)$ to be equivalent if their equivalence can be proved in T. Theorem 8 asserts that the equivalence class of a formula φ is determined by its semantics: that is, by how it is interpreted in the various models of M. Motivated by this observation, we ask the following:

Question 9. Let T be a first-order theory. Suppose that, for every model M of T, we specify a subset $\widetilde{M} \subseteq M$. Under what conditions does there exist a formula $\varphi(x)$ in the language of T such that $\widetilde{M} = M[\varphi]$ for every model M?

Of course, one should not expect such a formula to exist if the subsets $\widetilde{M} \subseteq M$ are chosen at random. At the very least, we need to demand that the subsets $\widetilde{M}$ have some relation to one another as the model M varies. To formulate this more precisely, we recall the following:

Definition 10. Let T be a first-order theory and let M and M' be models of T. A function $f : M \to M'$ is called an *elementary embedding* if for every first-order formula $\varphi(x_1, \ldots, x_n)$ in the language of T and every n-tuple of elements $a_1, a_2, \ldots, a_n$, we have

$$(M \vDash \varphi(a_1, \ldots, a_n)) \Leftrightarrow (M' \vDash \varphi(f(a_1), \ldots, f(a_n))).$$

Armed with this definition, we can state an obvious necessary condition in the setting of Question 9:

$(*)$ For every elementary embedding $f : M \to M'$ of models of T, we have $\widetilde{M} = f^{-1}(\widetilde{M'})$.

It follows immediately from the definition that condition $(*)$ is *necessary* for the collection of subsets $\widetilde{M} \subseteq M$ to be defined by some first-order formula $\varphi(x)$. However, it is not sufficient:

Example 11. Let T be the first-order theory of fields, so that the models of T are fields. For every field F, define

$$\widetilde{F} = \{x \in F : \ x \text{ is a root of unity } \} \subseteq F.$$

Then this collection of subsets satisfies condition $(*)$: in fact, for *any* field extension $F \to F'$, an element $a \in F$ is a root of unity in F if and only if its image is a root of unity in F'. However, there does not exist a single first-order formula $\varphi(x)$ such that $\widetilde{F} = \{a \in F : \varphi(a) \text{ is true in } F \ \}$.

We can prove this by appealing to the language of *ultraproducts*. Let $\mathbf{C}$ denote the field of complex numbers, let $\mathcal{U}$ be some nonprincipal ultrafilter on the set of positive integers, and let

$$F = (\prod_{n>0} \mathbf{C})/\mathcal{U}$$

be the ultraproduct of countably many copies of $\mathbf{C}$ determined by $\mathcal{U}$. The elements of F can be identified with equivalence classes of sequences of complex numbers $(z_1, z_2, z_3, \ldots)$, where we consider two sequences to be equivalent if they agree almost everywhere (in the sense of the ultrafilter $\mathcal{U}$).

A basic foundational result in model theory asserts that, for any first-order formula $\varphi(x)$, we have

$$(F \vDash \varphi(\{z_n\})) \Leftrightarrow (\mathbf{C} \vDash \varphi(z_n) \text{ for almost all } n\ ).$$

Suppose that $\varphi(x)$ was a formula that was satisfied precisely by the roots of unity in any field. Let $z \in F$ be the equivalence class of the sequence $(e^{2\pi i/1}, e^{2\pi i/2}, e^{2\pi i/3}, \ldots)$. Then $\mathbf{C} \vDash \varphi(e^{2\pi i/n})$ for every n, so $F \vDash \varphi(z)$. It follows that z is a root of unity in F: that is, $z^k = 1$ for some fixed integer k. Applying the preceding statement again, this time for the first order statement $x^k = 1$, we conclude that $e^{2\pi i/n}$ is a kth root of unity for almost all n. This contradicts our assumption that $\mathcal{U}$ is nonprincipal.

Makkai proved that this is essentially the only obstruction:

Theorem 12 (Makkai Definability Theorem)**.** *Let T be a first-order theory and suppose we are given a subset $\widetilde{M}$ for each model of M. Assume that:*

- $(*)$ *For every elementary embedding $f : M \to M'$ of models of T, we have $\widetilde{M} = f^{-1}(\widetilde{M'})$.*
- $(*')$ *For every collection of models $\{M_j\}_{j \in J}$ and every ultrafilter $\mathcal{U}$ on J with ultraproduct $M = (\prod_{j \in J} M_j)/\mathcal{U}$, we have*
$$(\{a_j\}_{j \in J} \in \widetilde{M}) \Leftrightarrow (a_j \in \widetilde{M_j} \text{ for almost all } j \in J).$$

Then there exists a formula $\varphi(x)$ in the language of T such that $\widetilde{M} = M[\varphi]$ for each model M of T.

Theorem 12 can be regarded as a natural companion to Gödel's completeness theorem: it describes the conditions under which there *exists* a formula $\varphi(x)$ with a certain semantic interpretation, while Gödel's theorem guarantees that $\varphi(x)$ is unique up to provable equivalence. Together, they represent a further step towards recovering the syntax of a theory T from its semantics. Our ultimate goal in this course will be to prove a further refinement of Theorem 12, also due to Makkai, which can be stated in a rough form as follows:

Theorem 13 (Strong Conceptual Completeness)**.** *Let T be a first-order theory and let* $\mathrm{Mod}(T)$ *be the category of models of T (where the morphisms are given by elementary extensions). Then T can be recovered, up to equivalence, from the structure of* $\mathrm{Mod}(T)$ *as a "category with ultraproducts."*

To formulate Theorem 13 precisely, we need to say what it should mean for two first-order theories to be considered equivalent. There are some obvious situations in which we might want to say this. For example, if T and T' share the same language, each axiom of T can be proven from the axioms of T', and each axiom of T' can be proven from the axioms of T, then we should consider T and T' to be equivalent. However, there are also more subtle forms that equivalence can take.

Example 14 (Projective Geometry)**.** Let k be a field and let $\mathbf{P}^2(k)$ be the projective plane over k. The *points* of $\mathbf{P}^2(k)$ are 1-dimensional subspaces $V \subseteq k^3$. We can also speak of *lines* in $\mathbf{P}^2(k)$, which are determined by two-dimensional subspaces $W \subseteq k^3$. We say that a point (given by a 1-dimensional subspace V) is incidence to a line (given by a 2-dimensional subspace W) if $V \subseteq W$.

More generally, an *abstract projective plane* consists of a set P of *points*, a set L of *lines*, and a binary relation between points and lines called *incidence*, which we write as $p \in \ell$. This relation is required to satisfy the flllowing axioms:

- $(A1)$ For every pair of distinct points $p, p' \in P$, there is a unique line $\ell \in L$ with $p \in L$ and $p' \in L$; we denote ℓ by $\overline{pp'}$.
- $(A2)$ For every pair of distinct lines $\ell, \ell' \in L$, there is a unique point $p \in P$ such that $p \in \ell$ and $p \in \ell'$; we denote p by $\ell \cap \ell'$.

In the case where $(P, L, \in)$ comes from a field k, we can recover k. More precisely, we first need to choose four points $a, b, c, d \in P$ which are *in general position* (no three of which are collinear), and we can then try to "coordinatize the plane" to recover a field k. This construction does not always work: for a general abstract projective plane $(P, L, \in)$, we need to assume that it satisfies *Pappus's hexagon theorem*:

(*A*3) Given two triples of collinear points $p, q, r \in P$ and $p', q', r' \in P$, the intersection points $\overline{pq'} \cap \overline{p'q}$, $\overline{pr'} \cap \overline{rp'}$, $\overline{qr'} \cap \overline{rq'}$ are also collinear.

However, this is all one needs. More precisely, there are mutually inverse constructions which pass back and forth between the datum of a field k, and the datum of an abstract projective plane $(P, L, \in)$ with four distinguished points in general position, satisfying Pappus's theorem.

Each of these data can be understood as a model for a suitably chosen first-order theory, but these theories look *a priori* very different. The theory of fields is usually formulated in terms of arithmetic notions (like addition and multiplication), while the theory of abstract projective planes is formulated in terms of points, lines, and incidence.

The focus of this course will be on the "categorical" approach to first-order logic, which emerges naturally from the problem of formulating a notion of equivalence between first-order theories that can accommodate situations like Example 14. One of our first goals will be to associate to each first-order theory T a category $\mathrm{Syn}(T)$ which we will refer to as *the syntactic category* of T. We will say that two first-order theories T and T' are equivalent if $\mathrm{Syn}(T)$ and $\mathrm{Syn}(T')$ are equivalent as categories. Theorem 13 can be stated more precisely as follows:

Theorem 15 (Makkai)**.** *Let T be a first-order theory. Then the syntactic category $\mathrm{Syn}(T)$ can be recovered from the category of models $\mathrm{Mod}(T)$. More precisely, there is a canonical equivalence*

$$\mathrm{Syn}(T) \simeq \mathrm{UFun}(\mathrm{Mod}(T), \mathcal{S}\mathrm{et})$$

where $\mathrm{UFun}(\mathrm{Mod}(T), \mathcal{S}\mathrm{et})$ denotes the category of "ultrafunctors" from $\mathrm{Mod}(T)$ to $\mathcal{S}\mathrm{et}$: that is, functors which are compatible with ultraproducts.

Remark 16. One can also go the other way: the category $\mathrm{Mod}(T)$ can be realized as a full subcategory of the category $\mathrm{Fun}(\mathrm{Syn}(T), \mathcal{S}\mathrm{et})$ of functors from $\mathrm{Syn}(T)$ to the category of sets. However, this is closer to being a definition than a theorem.

The syntactic category $\mathrm{Syn}(T)$ of a first-order theory T is an example of a very special kind of category, called a *pretopos*. Moreover, one can show that every pretopos arises in this way, provided that one adopts a sufficiently liberal interpretation of the phrase "first-order theory." More precisely, one needs to replace work with a nonclassical variant of first order logic, where one abandons the law of the excluded middle. The motivation for considering such generalizations need not come from commitment to any particular philosophical position: they can be motivated by pragmatic mathematical concerns.

Example 17. Let T be the theory of abelian groups described in Example 5. Then $\mathrm{Mod}(T)$ is a category whose objects are abelian groups, and whose morphisms $f : M \to N$ are *elementary embeddings* of abelian groups. Every elementary embedding is a group homomorphism, but not conversely. An elementary embedding $f : M \to N$ is required to preserve the truth of *any* statement which can be formulated in first-order logic. So, for example, if x is an element of M which is not divisible by 2, then $f(x)$ cannot be divisible by 2 in N. An arbitrary group homomorphism need not have this property.

On the other hand, there is a "nonclassical" theory T' whose category of models $\mathrm{Mod}(T')$ is the category whose objects are abelian groups and whose morphisms are group homomorphisms $f : M \to N$. These homomorphisms are required to preserve the truth of certain kinds of statements (such as equations $x+y = z$), but not others (such as inequations $x + y \neq z$). Passing from the setting of classical first order-logic to the setting of pretopoi accommodates such examples in a natural way.

As the terminology suggests, the notion of a pretopos is quite closely related to the notion of a *topos*, which is the other main subject of this course.

Remark 18. Throughout this course, we will always use the word *topos* to refer to what is known as a *Grothendieck topos*. There is also a more general notion of *elementary topos*, which we will not discuss.

Topos theory was introduced by Grothendieck and his students, motivated by problems in algebraic geometry. If X is an algebraic variety defined over the field $\mathbf{C}$ of complex numbers, then the set $X(\mathbf{C})$ of $\mathbf{C}$-valued points of X has a natural topology, which is induced by the usual topology on $\mathbf{C}$. This topological space is reasonably well-behaved (for example, if the algebraic variety X is smooth, then $X(\mathbf{C})$ is a manifold), and one can use it to define invariants (like homology and cohomology) which are very useful in the study of X.

Grothendieck and his students were interested in defining these sorts of invariants for algebraic varieties over other fields which have no interesting topology of their own (such as finite fields). In this case, one can still extract a topological space by endowing X with the *Zariski topology*: that is, by declaring a subset of X to be closed if it can be defined by polynomial equations. This topology is useful for some purposes, but not for extracting algebro-topological invariants like cohomology. As a remedy, Grothendieck introduced the notion of a *topos* and showed that the theory of (sheaf) cohomology could be extended from the setting of topological spaces to the setting of topoi. Moreover, he associated to each algebraic variety a particular topos (the *étale topos*) for which the resulting invariants turned out to be very powerful.

Topos theory is essentially a generalization of point-set topology. More precisely, there is a functor which associates a topos to every topological space, and that functor is fully faithful when restricted to the class of *sober* topological spaces (which includes all Hausdorff spaces). However, it is far from being essentially surjective: there are many interesting examples of topoi which do not come from topology. For example, every first-order theory T has a *classifying topos*, from which we can recover T up to equivalence (in the sense described above). So one can also think of topos theory as a generalization of first-order logic.

Remark 19. One can think of point-set topology and first-order logic as the study of two special classes of topoi. Moreover, these classes have nonempty intersection: the theory of Boolean algebras has an incarnation both in mathematical logic (as the theory of propositional logic) and in point-set topology (as the theory of *Stone spaces*: that is, compact Hausdorff spaces which are totally disconnected, like the Cantor set).

The relevant circle of ideas can be roughly pictured as follows:

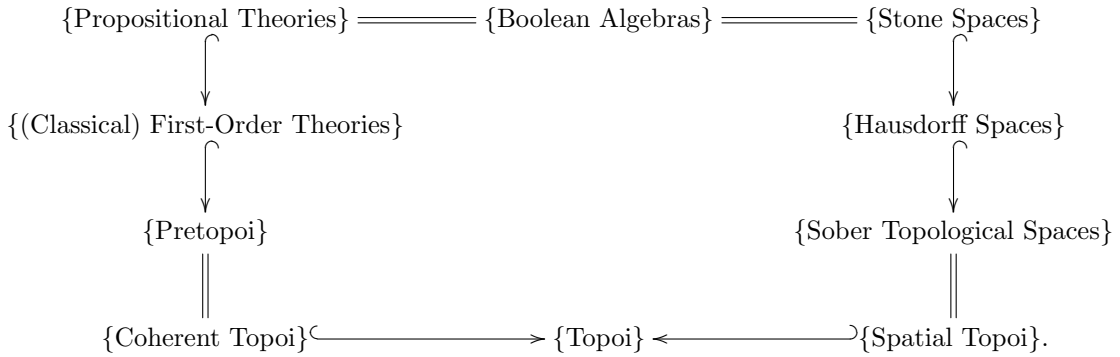

The trajectory of this course will take us counterclockwise around this diagram, beginning with first-order logic and ending with a proof of Theorem 15.