The Yoneda Lemma: What's It All About?

Tom Leinster

Department of Pure Mathematics, University of Cambridge
Email: leinster@dpmms.cam.ac.uk
Web: http://www.dpmms.cam.ac.uk/∼leinster

October 2000

1 The Yoneda Lemma **1**

2 Corollaries **6**
2.1 A representation is a universal element 6
2.2 The Yoneda embedding 8
2.3 Isomorphic representables 9

The level of generality and abstraction in the Yoneda Lemma means that many people find it quite bewildering. This document is meant to guide you slowly through what the Yoneda Lemma and its corollaries say, and give you some wider conceptual perspective. It contains no results other than those in the lectures, and as such is not 'required reading' for the course, but it might make your life easier.

Breathe deeply, take it slowly, and remain calm.

1 The Yoneda Lemma

Here's the statement of the Lemma. The proof was in lectures so I won't reproduce it here; in any case, once you have thoroughly understood the statement, you should find the proof straightforward.

The Yoneda Lemma *Let $\mathcal{C}$ be a locally small category. Then*

$$[\mathcal{C}^{\mathrm{op}}, \mathbf{Set}](H_A, X) \cong X(A) \qquad (*)$$

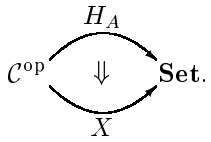
naturally in $A \in \mathcal{C}$ and $X \in [\mathcal{C}^{\mathrm{op}}, \mathbf{Set}]$.

First I will go through what this *says* at the formal level. Then I will try to explain what it *means* in a more intuitive sense.

Reminder: $(*)$ *says* $[\mathcal{C}^{\mathrm{op}}, \mathbf{Set}](H_A, X) \cong X(A)$

What it says Experience shows that many students are confused by the left-hand side of equation $(*)$. Let's dissect it.

- $\mathcal{C}$ is a category.
- $\mathcal{C}^{\mathrm{op}}$ is also a category, the opposite or dual of $\mathcal{C}$, obtained by keeping the same objects and reversing all the arrows.
- **Set** is a category too, whose objects are sets and whose morphisms are functions.
- For any two categories $\mathcal{A}$ and $\mathcal{B}$ there is a category $[\mathcal{A}, \mathcal{B}]$, whose objects are functors from $\mathcal{A}$ to $\mathcal{B}$ and whose morphisms are natural transformations.
- In particular we have the category $[\mathcal{C}^{\mathrm{op}}, \mathbf{Set}]$, whose objects are functors $\mathcal{C}^{\mathrm{op}} \longrightarrow \mathbf{Set}$ and whose morphisms are natural transformations.
- For any object A of $\mathcal{C}$, there is a functor $H_A : \mathcal{C}^{\mathrm{op}} \longrightarrow \mathbf{Set}$ (also written $\mathcal{C}(-, A)$). This functor is defined on objects by $H_A(B) = \mathcal{C}(B, A)$, and on morphisms by $H_A(f) = f^*$ (compose with f). (Think of dual vector spaces if it helps.)
- X is a functor $\mathcal{C}^{\mathrm{op}} \longrightarrow \mathbf{Set}$.
- For any category $\mathcal{D}$ and objects D, D' of $\mathcal{D}$, the set of morphisms from D to D' is written $\mathcal{D}(D, D')$.
- In particular, the left-hand side of $(*)$ is the set of morphisms in $[\mathcal{C}^{\mathrm{op}}, \mathbf{Set}]$ from H_A to X. That is, it is the set of natural transformations of the form



So the left-hand side of $(*)$ is a set. The right-hand side is also a set. Hence the isomorphism $(*)$ is a bijection between sets.

In summary, Yoneda says that a transformation from H_A to X is the same thing as an element of $X(A)$.

Digression: size worries If you are happy with this explanation then so much the better. But it does contain a slight economy with the truth: namely, that for a locally small category $\mathcal{C}$, the functor category $[\mathcal{C}^{\mathrm{op}}, \mathbf{Set}]$ is not in general locally small, and so the left-hand side of $(*)$ is *a priori* a class and not necessarily a set. However, when we prove the Yoneda Lemma we set up a bijection between this class and the right-hand side of $(*)$, which certainly *is* a set: hence the left-hand side is a set too.

It's really best not to worry about this kind of point if you can help it. For those who remain concerned, take $\mathcal{C}$ to be small rather than just locally small: this guarantees that $[\mathcal{C}^{\mathrm{op}}, \mathbf{Set}]$ is locally small, and your worries are over.

What it says, continued Putting these worries aside, we have seen that for fixed A and X, the Yoneda Lemma claims there is a bijection between a certain pair of sets. What about 'naturally in A and X'? Recall that if $F, G : \mathcal{D} \longrightarrow \mathcal{E}$ are a pair of functors, we use the phrase '$F(D) \cong G(D)$ naturally in $D \in \mathcal{D}$' to mean that there is a natural isomorphism $F \cong G$. The use of this phrase in the Yoneda Lemma carries the implication that each side of $(*)$ is functorial in both A and X; this means, for instance, that a map $X \longrightarrow X'$ induces a map

$$[\mathcal{C}^{\mathrm{op}}, \mathbf{Set}](H_A, X) \longrightarrow [\mathcal{C}^{\mathrm{op}}, \mathbf{Set}](H_A, X')$$

(as can be seen), and that there's a way of choosing the isomorphisms $(*)$ for *all* A and X which is compatible with such induced maps. So more exactly, what the Yoneda Lemma says is that the 'evaluation' functor

$$\begin{array}{ccc} \mathcal{C}^{\mathrm{op}} \times [\mathcal{C}^{\mathrm{op}}, \mathbf{Set}] & \xrightarrow{\mathrm{ev}} & \mathbf{Set} \\ (A, X) & \longmapsto & X(A) \end{array}$$

is naturally isomorphic to the composite functor

$$\mathcal{C}^{\mathrm{op}} \times [\mathcal{C}^{\mathrm{op}}, \mathbf{Set}] \xrightarrow{H_\bullet \times 1} [\mathcal{C}^{\mathrm{op}}, \mathbf{Set}]^{\mathrm{op}} \times [\mathcal{C}^{\mathrm{op}}, \mathbf{Set}] \xrightarrow{\mathrm{Hom}} \mathbf{Set}.$$

Here

$$H_\bullet : \mathcal{C} \longrightarrow [\mathcal{C}^{\mathrm{op}}, \mathbf{Set}]$$

is the 'Yoneda embedding', as detailed in lectures, which sends an object $A \in \mathcal{C}$ to the functor H_A and a morphism $f : A \longrightarrow A'$ to the natural transformation $H_f : H_A \longrightarrow H_{A'}$.

I will now suggest some ways of understanding the Yoneda Lemma.

Smaller formulae At a very practical level, you can think of the Yoneda Lemma as a useful tool: later in the course we'll come across various large and perhaps mystifying expressions, and by applying the isomorphism $(*)$ from left to right we will be able to reduce them to something more friendly.

Topological presheaves You can try to get a handle on the Yoneda Lemma by considering the following special case.

Fix a topological space S, and denote by $\mathcal{O}(S)$ the poset of open subsets of S, ordered by inclusion. Posets can be regarded as categories; thus an object of the corresponding category (also denoted $\mathcal{O}(S)$) is an open subset of S, and

$$\mathrm{Hom}(V, U) = \begin{cases} 1 & \text{if } V \subseteq U \\ \emptyset & \text{otherwise.} \end{cases} \qquad (\dagger)$$

The functor category $[\mathcal{O}(S)^{\mathrm{op}}, \mathbf{Set}]$ is called the category of *presheaves* on S. (This year's algebraic geometry course uses a different way of defining (pre)sheaves; the two approaches are equivalent, but don't worry about this here.) Explicitly, a presheaf X on S consists of

- for each open $U \subseteq S$, a set $X(U)$
- for each open V and U with $V \subseteq U$, a function $X(U) \longrightarrow X(V)$, usually written $p \longmapsto p|_V$ and called 'restriction'.

Restriction is required to satisfy functoriality axioms: $(p|_V)|_W = p|_W$ and $p|_U = p$, for $p \in X(U)$ and $W \subseteq V \subseteq U$. The classic example of a presheaf on a space is where $X(U)$ is the set of continuous functions from U to the real numbers, and restriction is restriction in the usual sense. A *morphism* $\alpha : X \longrightarrow Y$ of presheaves is a natural transformation, and explicitly consists of a family $(\alpha_U : X(U) \longrightarrow Y(U))_{U \in \mathcal{O}(S)}$ of functions satisfying $(\alpha_U(p))|_V = \alpha_V(p|_V)$ for each $p \in X(U)$ and open $V \subseteq U$.

The *representable* presheaves are those of the form $H_U : \mathcal{O}(S)^{\mathrm{op}} \longrightarrow \mathbf{Set}$, where $U \subseteq S$ is open. Then $H_U(V)$ is given by the formula (†), and the restriction maps for H_U are uniquely determined.

Now ask yourself: given an open $U \subseteq S$ and a presheaf X on S, what's a morphism $H_U \longrightarrow X$? Well, it's a family $\alpha_V : H_U(V) \longrightarrow X(V)$ of functions, one for each open V, which is compatible with the restriction maps. After some contemplation you should see that such an α is entirely determined by the value of α_U at the single element of $H_U(U)$. So a map $H_U \longrightarrow X$ is just the same thing as an element of $X(U)$: that is, there is a bijection

$$[\mathcal{O}(S)^{\mathrm{op}}, \mathbf{Set}](H_U, X) \cong X(U).$$

And of course, this is the Yoneda Lemma (minus naturality) in the case $\mathcal{C} = \mathcal{O}(S)$.

Monoid actions Here is another potentially enlightening special case.

Fix a monoid M. As we have seen, monoids are the same thing as one-object small categories, and viewing M in this way, $[M^{\mathrm{op}}, \mathbf{Set}]$ is the category of right M-sets (= sets equipped with a right action by M). If we write A for the single object of the category M, then the representable functor $H_A : M^{\mathrm{op}} \longrightarrow \mathbf{Set}$ corresponds to what is sometimes called the 'right regular representation of M': that is, the set M acting on itself by composition. I will (perhaps confusingly) write M for this particular right M-set. Then, for an arbitrary right M-set X, a morphism $\alpha : M \longrightarrow X$ of M-sets is entirely determined by $\alpha(1)$. Hence

$$[M^{\mathrm{op}}, \mathbf{Set}](M, X) \cong X,$$

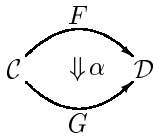
in accordance with Yoneda again.

Coherence General category theory springs no nasty surprises: any sensible equation you can write down is true. People sometimes say 'all diagrams commute'. Of course, you need to take this with a pinch of salt (and it's no excuse for omitting the proper checks in an exam...). But contrast, for instance, group theory, where there are plenty of equations, such as $a \cdot b = b \cdot a$, which are perfectly sensible but in general false. In category theory this equation wouldn't

in general make sense: if a composite $a \circ b$ exists then the composite $b \circ a$ usually doesn't.

To put it another way, in general category theory, there's at most one way of taking inputs of given types and obtaining an output of a given type. More snappily, there's only one way of getting from A to B. For example:

- Given a natural transformation



  and a map $C \xrightarrow{f} C'$ in $\mathcal{C}$, there's precisely one way of obtaining a map $FC \longrightarrow GC'$ in $\mathcal{D}$: naturality of α says that the two routes from top-left to bottom-right in the square

$$\begin{array}{ccc} FC & \xrightarrow{Ff} & FC' \\ {\scriptstyle \alpha_C}\downarrow & & \downarrow{\scriptstyle \alpha_{C'}} \\ GC & \xrightarrow[Gf]{} & GC' \end{array}$$

  are equal.

- Given a functor $F : \mathcal{C} \longrightarrow \mathcal{D}$ and maps $A \xrightarrow{f} B \xrightarrow{g} C$ in $\mathcal{C}$, there's only one way of building a map $FA \longrightarrow FC$: functoriality implies that the two maps $F(g \circ f)$ and $F(g) \circ F(f)$ are actually the same.

- For 'general' sets A, B, C, there's only one sensible isomorphism

$$(A \times B) \times C \xrightarrow{\sim} A \times (B \times C).$$

The Yoneda Lemma is another example. We have two ways of taking as input a pair $(A \in \mathcal{C}, X : \mathcal{C}^{\mathrm{op}} \longrightarrow \mathbf{Set})$, and producing as output a set: one appears as the left-hand side of $(*)$, and the other as the right. Yoneda says that these two ways are, in fact, the same.

If Yoneda *weren't* true then the world would look very different, and much more complex. Starting simply from a functor $X : \mathcal{C}^{\mathrm{op}} \longrightarrow \mathbf{Set}$, we would obtain a new functor

$$X' = [\mathcal{C}^{\mathrm{op}}, \mathbf{Set}](H_{\bullet}, X) : \mathcal{C}^{\mathrm{op}} \longrightarrow \mathbf{Set},$$

and hence a whole sequence of functors $X, X', X'', X''', \ldots$, potentially all different. In reality they are all the same.

2 Corollaries

The Yoneda Lemma has (at least) three corollaries. Each can be proved directly as well, and in fact that's not a bad exercise.

2.1 A representation is a universal element

Corollary *Let $\mathcal{C}$ be a locally small category and $X : \mathcal{C}^{\mathrm{op}} \longrightarrow \mathbf{Set}$. Then a representation of X consists of an object A of $\mathcal{C}$ together with an element $u \in X(A)$ such that*

$$\begin{array}{l} \textit{for any } B \in \mathcal{C} \textit{ and } x \in X(B), \textit{ there is a unique map} \\ f : B \longrightarrow A \textit{ satisfying } (Xf)(u) = x. \end{array} \qquad (\ddagger)$$

To clarify the statement, first recall that a representation of X is, by definition, an object A of $\mathcal{C}$ together with a natural isomorphism $\alpha : H_A \longrightarrow X$. The Corollary says that such pairs (A, α) are in one-to-one correspondence with pairs (A, u) satisfying $(\ddagger)$. This follows easily from the Yoneda Lemma.

We can think of u as a 'universal' or 'generic' element. I will try to explain what's going on by two examples.

Example 1 Fix vector spaces U and V, and consider the functor

$$\begin{array}{rccl} \mathrm{Bilin}(U, V; -) : & \mathbf{Vect} & \longrightarrow & \mathbf{Set}, \\ & W & \longmapsto & \mathrm{Bilin}(U, V; W) \\ & & & = \{\text{bilinear maps } U \times V \longrightarrow W\}. \end{array}$$

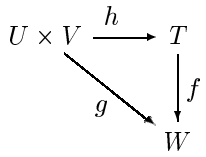
Then a representation of $\mathrm{Bilin}(U, V; -)$ can be described in either of two equivalent ways:

a. as a vector space T together with an isomorphism

$$\mathbf{Vect}(T, W) \cong \mathrm{Bilin}(U, V; W)$$

natural in $W \in \mathbf{Vect}$

b. as a vector space T together with a bilinear map $h : U \times V \longrightarrow T$, such that

> for any vector space W and bilinear $g : U \times V \longrightarrow W$, there is a unique linear $f : T \longrightarrow W$ making
>
> 
>
>
> commute.

Part (a) is just the definition of representation. Part (b) is the description given in the Corollary, or rather the dual of the Corollary (concerning *co*variant functors $X : \mathcal{C} \longrightarrow \mathbf{Set}$: try writing out this statement). The map called h should for consistency be called u (but would then look like an element of U). Those who know about such things will recognise T as the tensor product $U \otimes V$, and h as the map $(u, v) \longmapsto u \otimes v$.

You will observe that the first description is substantially shorter than the second. Indeed, it's clear enough that if the situation of (b) holds then there is certainly an isomorphism

$$\mathbf{Vect}(T, W) \longrightarrow \mathrm{Bilin}(U, V; W)$$

natural in W, got by composition with h. But it looks at first as if (b) says rather more than (a): that not only are the two things naturally isomorphic, they are naturally isomorphic in a rather specific manner. The Corollary tells us that this is an illusion: all such natural isomorphisms arise in this manner. It's the word 'natural' in (a) that hides all the explicit detail.

Example 2 Let $\mathcal{C} \underset{G}{\overset{F}{\underset{\longleftarrow}{\overset{\longrightarrow}{\bot}}}} \mathcal{D}$ be an adjunction, and fix an object A of $\mathcal{C}$. Then the functor

$$\mathcal{C}(A, G-) : \mathcal{D} \longrightarrow \mathbf{Set}$$

is representable, as can be expressed in either of the following two ways:

a. $\mathcal{C}(A, GB) \cong \mathcal{D}(FA, B)$ naturally in $B \in \mathcal{D}$

b. the unit map $\eta_A : A \longrightarrow G(FA)$ is an initial object of the comma category $(A \Rightarrow G)$.

Again, the first description comes from the definition of representability, and the second from (the dual of) the Corollary. (It takes a moment to see this; I leave that to you.) We looked at the second description from another perspective in section B of lectures.

2.2 The Yoneda embedding

The next result is a corollary of the Yoneda Lemma (and not of the previous corollary).

Corollary *For any locally small category $\mathcal{C}$, the functor*

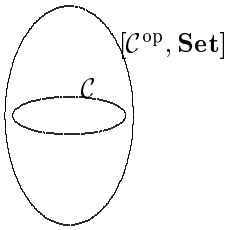
$$H_\bullet : \mathcal{C} \longrightarrow [\mathcal{C}^{\mathrm{op}}, \mathbf{Set}]$$

is full and faithful.

The functor $H_\bullet$ is known as the *Yoneda embedding*. The fact that it's faithful is what justifies the name 'embedding'; the fact that it's also full means it's an especially nice kind of embedding—a map $H_A \longrightarrow H_B$ in $[\mathcal{C}^{\mathrm{op}}, \mathbf{Set}]$ is exactly the same as a map $A \longrightarrow B$ in $\mathcal{C}$.

It's also true that $H_\bullet$ is injective on objects (if $A \neq B$ then $H_A \neq H_B$), which follows from our convention that the hom-sets of a category are disjoint (Remark A1.3(a) of the notes). But this should be regarded as unimportant: we really aren't interested in equality of objects in a category, only isomorphism.

Anyway, the fact that $H_\bullet$ is full and faithful (and injective on objects) means that we can regard $\mathcal{C}$ as sitting inside $[\mathcal{C}^{\mathrm{op}}, \mathbf{Set}]$ as a full subcategory:



Later in the course we'll see how any functor $\mathcal{C}^{\mathrm{op}} \longrightarrow \mathbf{Set}$ can be built out of representables H_A, in very roughly the same way that any number is built as a product of primes.

2.3 Isomorphic representables

The previous corollary has in turn the following corollary:

Corollary *For objects A and B of a locally small category $\mathcal{C}$,*

$$H_A \cong H_B \iff A \cong B \iff H^A \cong H^B.$$

The force of this is that $H_A \cong H_B \Rightarrow A \cong B$; the other direction of implication follows immediately from functoriality of $H_\bullet$. The second $\iff$ is just the dual result, and therefore also follows immediately.

The Corollary can be explained as follows. Regard $H_A(U) = \mathcal{C}(U, A)$ as 'A viewed from U': then our result says that two objects are the same if and only if they look the same from all viewpoints.

The category of sets is very unusual in this context: for sets A and B,

$$A \cong B \iff H_A(1) \cong H_B(1),$$

and so the Corollary has a trivial proof for $\mathcal{C} = \mathbf{Set}$. In other words, in **Set** it's enough to look at everything from the one-element set 1—the only thing that matters about a set is its elements!

In contrast, take $\mathcal{C} = \mathbf{Gp}$. Imagine that we have two groups A and B, and someone is telling us that A and B 'look the same from U' for various groups U. Then, for instance,

- $H_A(1) \cong H_B(1)$ would tell us nothing at all
- $H_A(\mathbb{Z}) \cong H_B(\mathbb{Z})$ would tell us that A and B have isomorphic underlying sets—that is, the same cardinality, but perhaps quite different group structures
- $H_A(\mathbb{Z}/p\mathbb{Z}) \cong H_B(\mathbb{Z}/p\mathbb{Z})$ would tell us that A and B have the same number of elements of order p, for a prime p,

and so on. Each of these only gives partial information about the similarity of A and B, but the whole natural isomorphism $H_A \cong H_B$ tells us that $A \cong B$.

