

1. Introduction

When you first come to university, and study math, you start out with looking at matrices. They let you package systems of linear equations into a single matrix one. And in the beginning, you learn about Gaussian elimination. It's similar to how one would normally treat simultaneous equations in school - you can implement addition and multiplication of equations as row operations on the corresponding matrix, and reduce it to some canonical form, often called "row echelon form" or "reduced row echelon form".

And later on, you might learn a neat fact that relates these row operations to more standard matrix multiplication. Any row operation on some matrix A can be implemented by just multiplying on the left with a fixed matrix M - and the way you find M is just by applying the operation to the identity matrix.

For example, if we take $\begin{pmatrix} 1 & 2 \\ 5 & 7 \end{pmatrix}$, and we want to implement the operation of subtracting 5 times the first row from the second row, to leave us with $\begin{pmatrix} 1 & 2 \\ 0 & -3 \end{pmatrix}$, then we can use the following procedure. Apply the same operation to the identity matrix, so $\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ goes to $\begin{pmatrix} 1 & 0 \\ -5 & 1 \end{pmatrix}$, and one can verify that $\begin{pmatrix} 1 & 0 \\ -5 & 1 \end{pmatrix} \begin{pmatrix} 1 & 2 \\ 5 & 7 \end{pmatrix} = \begin{pmatrix} 1 & 2 \\ 0 & -3 \end{pmatrix}$. So, when we solve a linear system $A\vec{x} = \vec{b}$ using row operations, what we're really doing is multiplying on the left by a series of matrices, $R_1, R_2, \dots$, until we reduce the system to something more tractable (e.g. a triangular matrix)!

But - why is this true? It turns out this is an instance of a much more general fact called the Yoneda Lemma, said to be the


fundamental theorem of category theory. (Write category theory). Now, I don't have time to go through all the definitions required to state the Yoneda Lemma, and what category theory is - instead, I'm gonna cut to the chase and tell you what category theory *does*, and why you should care about it. To understand, we need to understand one of its fundamental concepts - *naturality*.

2. What is Naturality?

2.1 Invariance


Mathematicians often use words like "natural" and "canonical", and one of the major successes of category theory was to nail down precisely what this might mean. But when you hear the word natural, the first thing that might come to mind is "natural number". So let's start with an... easier? question. What is a number?

Well, here's where I think numbers originate from. You have a bunch of different collections of things - 2 apples, 2 oranges, 2 doors, 2 cars, 2 hands. And you ask - what do they all have in common? What's the invariant?



Well, they all have 2 objects! We have a function, "size" or "cardinality", which assigns, to every collection, a number.

But there's something more interesting here - if we have a way of relating inputs, so them being in bijective correspondence, then we have a way of relating the resulting outputs - they're equal! This is the simplest kind of naturality - *invariance*.



Notice that “size” isn’t the *only* invariant we could’ve considered. Another invariant is being able to split the collection into equal halves - i.e., “evenness”. “Primeness” is also another invariant. There are even some cultures which don’t have names for numbers bigger than 5, and just have a single word for “many”- and that works as a system of invariants.

But there’s a sense in which numbers are special - every other invariant follows from them! So, if you want to study invariants of finite sets, you can decompose it as first determining their size, which we call “counting” or “combinatorics”, and then seeing how numbers determine all the other invariants, which we call “number theory”.

Ok so what’s next - I think most people can handle the concept of invariance. Well, after you learn to count, the next thing you probably learn is how to add! And that gets us to the second kind of naturality - *covariance*.

2.2 Covariance

Say you have 2 cherries and 3 bananas. You can combine them into a single collection of 5 fruits by taking their “union”, and then count how many there are, giving their size.





Figure 1: Numbers determine the invariants of finite sets.

Or, you can take their sizes individually, and then. . . add the numbers together! Indeed, this is basically what defines addition, and also gives us what we call covariance.




The idea is - we have a function, size, which relates inputs to outputs. That’s what all functions do. This function is “natural” if the following is true - whenever we have a way to relate inputs, in this case union, there’s a corresponding way to relate the outputs!

Multiplication is another example - in this case, we can take the cartesian product and make a grid of all possible choices, and then count them up, or we can take the sizes individually, and then multiply the numbers together!



The useful part is that both ways give you the same result, so you have the freedom to choose which one to perform. For example, counting votes. It's a lot easier to count votes in individual places and then add the numbers together, rather than bringing all the votes in the country to a single place and then count them up. With multiplication, it's almost always easier to multiply the numbers together than list out all possible choices.


A more complicated example might be power set - given any set X , you can associate to it the power set, $P(X)$. It's a function that relates inputs to outputs, so Y to $P(Y)$.




But moreover, it's natural - if you have a way to relate inputs, so a function $f : X \rightarrow Y$, then there's a corresponding way to relate the outputs - namely, given any subset of X , you can take the *image* under f to get a subset of Y . This is what covariance is about - upgrading a function between sets to a function between graphs. It's about recognising that your inputs and outputs aren't just separate objects, but are related in some ways. They're kind of like "graph homomorphisms".

2.3 Contravariance

In fact, the power set construction gives us an example of the third kind of naturality - contravariance! Namely, a way to relate X to Y , so a function $f : X \rightarrow Y$, gives us a way to relate $P(Y)$ to $P(X)$ - the preimage!




There's actually a nice description of this just in terms of function composition - you can identify subsets with their indicator functions, so you can think of $P(X)$ as the collection of functions from $X \rightarrow \{0, 1\}$. Then, the preimage is just function composition for indicator functions!



You also might remember this from school when having to work with graph transformations - you've got a function from x to y , and a change of coordinates x' . You have to reverse the change to get the arrow from x' to x - indeed, it's a general principle that functions are covariant in the output, but contravariant in the input.

There's also a simpler, more concrete example of contravariance - units! Say you're measuring a stick. Given a unit of measurement, you get back a numerical value - so, cm gives 50. And given a different one, you get a different value - m gives 0.5. There's a way to relate these inputs - a metre is 100 centimetres. And the way to relate the outputs "goes backwards" - 50 is 100 times 0.5.




Of course, here you can just view it as an example of covariance since the output arrow is reversible, whereas that's not true for say preimage. But I think it's useful to emphasise the "inverse" relation on outputs - in fact, the same reasoning is why vector components transform by P^{-1} . Imagine a function assigning, to each basis, the components of the vector in that basis. Moreover, what's often useful is that covariance and contravariance can combine to give invariance. This is why we always write down measurements with both the unit and the number - the unit changes covariantly, the number changes contravariantly, and together they give an invariant notion.

So that was a whistle-stop tour of naturality. Let's get back to matrices!

2.4 Matrices

It turns out it'll be a little easier to consider *column* operations first - we can always transpose our result to get the version for row operations.

So, a column operation takes in a matrix A with some fixed number of columns, and outputs another matrix $C(A)$. And it takes in another matrix B and outputs a matrix $C(B)$.



Now, my claim is that column operations are natural! So, if there's a way to relate A with B , specifically $B = MA$ (it's this way around because A and B have the same number of columns), then there's a way to relate the outputs, so $C(B) = MC(A)$.

Why? Well, each column of $C(A)$ is of the form $\sum \lambda_i A_i$ for A_i the i th column of A . Then, the j column of $MC(A)$ is $M \sum \lambda_{ij} A_i = \sum \lambda_{ij} M A_i = C(MA) = C(B)$. It's just linearity!

How does this help us? Well, now that we understand naturality, we're ready to employ a trick - we can write $A = AI$ for I the identity matrix. Then $C(A) = C(AI) = AC(I)$. And that gives our result - the column operation is just right-multiplication by some matrix, and we find it by applying it to the identity matrix!

We can also transpose everything and get the analogous results for row operations:

- Row operations are natural, so if $B = AM$, then $R(B) = R(AM) = R(A)M$.
- Then, you can write any matrix A as IA for I the appropriate identity matrix.
- So, $R(A) = R(IA) = R(I)A$. So you multiply on the left by a fixed matrix.

There's a few useful things we immediately get from this result: - It gives us a quick way to compose column operations - multiply the corresponding matrices!

Or, apply them in succession to the identity matrix. - By taking determinants, we see that any column operation only scales the determinant of the matrix by a constant factor, $\det(C(I))$, when acting on square matrices. - A column operation is invertible if and only if the corresponding matrix is invertible.

I want to emphasise that - the proof wasn't really hard, at all! The point was more creating a useful definition, *naturality*, that made the problem simple and transparent - which is often what category theory feels like. A big part of maths is the study of functions, things which relate inputs to outputs. Naturality asks you to consider the possibility that the inputs are related to each other - rather than being disparate,

solitary individuals, they form a kind of community, a graph, a network. The function is natural if, whenever there's a way to relate inputs, there's a corresponding way to relate outputs. It's a pattern in the way mathematicians seem to think and do mathematics.

Finally, I'd be remiss to not draw the connection to Cayley's theorem from group theory, which the Yoneda Lemma is often seen as a generalization of. We saw that every matrix M implements a row operation, $A \rightarrow MA$. This same map appears in Cayley's Theorem, where you replace a group element g with the map $k \rightarrow gk$, which satisfies a naturality condition similar to the one for matrices.

In general, what the Yoneda Lemma tells you is that what an object is is isomorphic to what an object does, how it acts on or relates to other objects.

And this is often a useful perspective to take - you can specify what you want your object to do, how you want to use it, and Yoneda will guarantee that this specifies the object basically uniquely. Often called a "universal property". In general, category theory cares a lot more about what an object does rather than what an object is, and Yoneda tells you these are equivalent.