FUNDAMENTAL CONCEPTS IN CALCULUS

DEANE YANG

1. Real number

A real number is assumed to be any number that can be written as a decimal with a finite number of digits to the left of the decimal point and possibly an infinite number to the right of the decimal.

But with the following caveat: If the tail of a decimal is an infinite sequence of 9's, then it is equal to the decimal, where the 9's are deleted and the last digit to the left of the 9's is incremented by 1. For example,

$$1.99999\cdots = 2$$

We will assume that we know how to do arithmetic (addition, subtraction, multiplication, division) with real numbers, even though it is tricky if there are an infinite number of digits.

2. Function

A function is a systematic process (the technical term is algorithm) for taking one or more inputs and generating one or more outputs. The functions studied in single variable calculus all take a single input, a real number, and produces a single output, also a real number.

Some simple examples are the following:

- A function, where, no matter what the input is, the output is 3. This is called a constant function.
- A function, where the output is always the same as the input unchanged.
- A function, where the output is the input squared.

In symbols, if the function is named f and the input is named x, then the output is written as $f(x)$.

In general, if you are handed a function named Ξ and a number named $\Diamond$, then you should always think of Ξ as a process (and not a number or a set of numbers), where if you feed it $\Diamond$ as an input, then the process will produce an output, which is written abstractly as $\Xi(\Diamond)$.

When you want to explain details of a function Ξ, you usually write this as

$$\Xi(\Diamond) = \text{ formula involving } \Diamond.$$

The functions above can therefore be described symbolically as follows:

- $h(z) = 3$: The function h, given any input z, produces the output 3.
- $p(u) = u$: The function p, given any input u, produces the output u
- $q(c) = c^2$: The function q, given any input c, produces the output c^2.

Note that the function f, defined by

$$f(t) = t^2 - 1$$

is the same function as the function h, defined by

$$h(y) = y^2 - 1,$$

because, given any input, h will produce the same output as f.

Once you have defined a function, say the function q above, and want to feed it a specific input, say *blah*, then the corresponding output is written as $q(blah)$. For this particular function, $q(blah)$ is equal to number consisting of *blah* squared. And $q(c + 2) = (c + 2)^2$. If a function f is defined to be $f(x) = x - 1$, then $f(x + 2) = (x + 2) - 1 = x + 1$.

3. Domain of a function

Often we want to restrict which real numbers are allowed to be an input to a particular function. The set of permissible inputs to a function f is called its domain. Sometimes, the domain is restricted, because the formula used to define f fails for some input values. Sometimes, however, we impose a restriction on the domain, even if the definition of the function is valid on a larger domain.

4. Continuity of a function

We can think of a function as a knob and a needle on a dial. The position of the knob represents the input, and the position of the needle represents the output. When the knob is adjusted, the needle need not move in the same direction as the knob. It also need not even move at all.

The rough idea of what it means for a function to be continuous is that, when we turn the knob, the needle never instantaneously jumps to a new position. This, however, turns out to be very tricky to describe in a precise way. We need to somehow distinguish between a needle instantaneously jumping to a new position, without actually traveling through the positions between the old and new position (this is actually impossible physically), and a needle moving from the old position to a new position very quickly.

Continuity is a necessary but not sufficient assumption for the discussion below. We will avoid using it directly but restricting our attention to functions that are continuous.

5. Derivative of a Function

It is natural to ask how sensitive the output of a function is to a change in the input. So we can define the sensitivity of a function to be the ratio

$$\text{Sensitivity of } f \simeq \frac{\text{Resulting change in output}}{\text{Change in input}}$$

For simple functions, this idea makes sense. For example, if the function is a constant function, where the output is always the same no matter what the input is, the sensitivity is clearly zero. If the output is always equal to the input, then the sensitivity is 1. If the output is always equal to twice the input, the sensitivity is 2.

For more complicated functions, the concept of sensitivity becomes more complicated and leads to the concept of the derivative of a function. The two main complications are the following:

- The sensitivity might depend on what input you start with
- The sensitivity might depend on how much you change the input

An example is the function $f(x) = x^2$.

Same starting input, different input changes:

Old input	Old output	New input	New output	Input change	Output change	Sensitivity
1	1	1.1	1.21	0.1	0.21	2.1
1	1	1.01	1.0201	0.01	0.0201	2.01

Different starting inputs, same input changes:

Old input	Old output	New input	New output	Input change	Output change	Sensitivity
1	1	1.1	1.21	0.1	0.21	2.1
2	4	2.1	4.41	0.1	0.41	4.1

The deep insight is that if we stay with the same starting input but calculate the sensitivity for smaller and smaller input changes, the sensitivity moves closer and closer to a single fixed value. We say that the sensitivity converges to a limiting value.

This can be illustrated with the same function $f(x) = x^2$. If we use 1 as the starting input and call the input change c, then we get the following:

Old input	Old output	New input	New output	Input change	Output change	Sensitivity
1	1	$1+c$	$(1+c)^2$	c	$2c + c^2$	$2+c$

If the change c is made smaller and smaller, the sensitivity gets closer and closer to the value 2. We call this limiting sensitivity the *derivative of f for the input* 1 and write it as $f'(1)$.

More generally, we can do this for any starting input x:

Old input	Old output	New input	New output	Input change	Output change	Sensitivity
x	x	$x + c$	$(x+c)^2$	c	$2xc + c^2$	$2x + c$

Again, by letting $c \to 0$, we see that the sensitivity converges to $2x$. Therefore, the derivative of f for the input x is $2x$. We write this as

$$f'(x) = 2x.$$

Notice that the derivative f' is itself also a function.

So what does this mean in practice? It means that for the function $f(x) = x^2$, if we start with an input value of, say, 3 and change the input a little, say 0.04, then the output will change by about twice the input, namely 0.08. Since $f(3) = 9$, this means that $f(3.04)$ is approximately $9 + 0.08 = 9.08$.

More generally, since

$$\text{Sensitivity} \simeq \frac{\text{Output change}}{\text{Input change}},$$

we can estimate the new output by

$$\text{New output} = \text{Old output} + \text{Output change}$$
$$\simeq \text{Old output} + (\text{Sensitivity})(\text{Input change}).$$

We say this using symbols as follows: Suppose for a function f and a starting input s, we know the output $f(s)$, and the derivative $f'(s)$. Then for another input value t close to s, we can estimate the output $f(t)$ by

$$(1) \qquad\qquad f(t) \simeq f(s) + f'(s)(t - s).$$

This is commonly known as the linear or tangent line approximation of f near s.

6. Integral

The integral of a function is designed to reconstruct a function from its derivative. The principle is that if you know how sensitive the output of a function is to small changes of the input, then you can get good estimates of how the output changes for a large change in input.

So suppose you know in advance the derivative f' of a function f but don't know f itself, except its value for one specific input s. We now want to find a way to estimate the value of $f(t)$ for some other input t. Using the linear approximation (1), we know how to estimate $f(t)$, if t is close to s. But what if t is not close to s?

The answer is quite simple: A big input change is just a sum of small input changes, so we can estimate the output change due to the big input change by adding up the output changes due to the small input changes.

For example, suppose we have a function f, where we know that $f'(x) = 2x$ (for all x) and $f(0) = 0$, and we want to find $f(1)$. We can use the linear approximation (1) to estimate, one

at a time, the values of $f(0.25), f(0.5), f(0.75), f(1)$. In other words, we just use repeatedly the estimate

$$\text{New output} \simeq \text{Old output} + (\text{Sensitivity})(\text{Input Change})$$

This gives

$$f(0.25) \simeq f(0) + f'(0)(0.25 - 0) = 0$$
$$f(0.5) \simeq f(0.25) + f'(0.25)(0.5 - 0.25) = 0.5(0.25) = 0.125$$
$$f(0.75) \simeq f(0.5) + f'(0.5)(0.75 - 0.5) = 0.125 + 1(0.25) = 0.375$$
$$f(1)) \simeq f(0.75) + f'(0.75)(1 - 0.75) = 0.375 + 1.5(0.25) = 0.75$$

This is too crude, but we could redo the calculation but using smaller incremental input changes. If we were to do this, we would find that the estimates for $f(1)$ get closer and closer to a specific value. For this specific example, it would turn out to be the value 1.

This process can be carried out using the derivative of any (not too weird) function f. It allows us to reconstruct the function f knowing only its value for one input and its derivative for all inputs.

In words, the difference between the outputs for two inputs can be estimated by summing a sequence of output changes due to a sequence of small input changes. Each output change due to a small input change can be estimated using the derivative (sensitivity) of the function. By making the small input changes approach zero, the estimate converges to a limiting value that is equal to the exact output change.

This process of estimating the output change using f' and taking a limit is encoded in the concept of the definite integral of f' from s to t, which is written

$$\int_s^t f'(x)\, dx.$$

The fact that the limiting value is the exact value of the output change is known as the Fundamental Theorem of Calculus, which says

$$f(t) - f(s) = \int_s^t f'(x)\, dx.$$

This process allows us to reconstruct the entire function f from knowing only its value for a single input and its derivative. If, for example, we know the value of f for only a single input t_0 and but its derivative $f'(t)$ for all inputs t, we can reconstruct, for any input t, the value of $f(t)$ using the Fundamental Theorem of Calculus,

$$f(t) = f(t_0) + \int_{t_0}^t f'(x)\, dx.$$

What happens if we integrate the function f itself? The process we applied above to the function f' can be applied to any function, even if the function is not known to be the derivative of some other function. What does

$$\int_s^t f(x)\, dx$$

mean? It means just do to f whatever we did to f' above. In practical situations, this integral does not always have any useful meaning. Mathematically, it allows us to construct a new function whose derivative is f. In other words, any function f is in fact the derivative of some other function g. We say that g is the antiderivative of f.

It is worth noting that, in most real life uses of calculus, the estimation process as described above, as well as more sophisticated versions of it, is more important than the mathematically precise definition and calculations of an integral.

7. POSITION AND VELOCITY

Suppose an object is moving back and forth along a straight line, and we're keeping track of its position $p(t)$ at time t. We can estimate the velocity of the object at any time by

$$\text{(2)} \qquad \text{Velocity} \simeq \frac{\text{Change in position}}{\text{Time elapsed}}.$$

This, however, is the same as what we earlier called the sensitivity of p. If we make the time interval shorter and shorter, the corresponding velocities (i.e., sensitivities) converge to what we call the instantaneous velocity of the object (i.e., the derivative of p). In other words, if v is the instantaneous velocity function, then $v = p'$.

This in turn means that if we know the velocity, then we can estimate the change in the object's position during a short time interval by

$$\text{(3)} \qquad \text{Change in position} \simeq (\text{Velocity})(\text{Time elapsed}).$$

If we know the instantaneous velocity of the object at every time during some time interval and the position of the object at the start of the time interval, then we can estimate the position of the object at the end of the time interval as follows:

(1) Divide the time interval into a sequence of small subintervals.
(2) Estimate the net change in position during each small subinterval using equation (3).
(3) Add up all the estimated position changes for the subintervals to get an estimated position change for the full time interval.

If we then redo this for smaller and smaller divisions of the time interval, the estimate converges to the exact change in position over the full time interval. Since this process is exactly the same as the one described in the previous section for calculating the integral of a derivative, it means that, if $v(t)$ is the instantaneous velocity at time t,

$$p(t_1) - f(t_0) = \int_{t_0}^{t_1} v(t)\, dt.$$