

Chapter udf

Functions

fun.1 Basics

sfr:fun:bas:
sec A *function* is a map which sends each **element** of a given set to a specific **element** in some (other) given set. For instance, the operation of adding 1 defines a function: each number n is mapped to a unique number $n + 1$. explanation

More generally, functions may take pairs, triples, etc., as inputs and returns some kind of output. Many functions are familiar to us from basic arithmetic. For instance, addition and multiplication are functions. They take in two numbers and return a third.

In this mathematical, abstract sense, a function is a *black box*: what matters is only what output is paired with what input, not the method for calculating the output.

Definition fun.1 (Function). A *function* $f: A \rightarrow B$ is a mapping of each **element** of A to an **element** of B .

We call A the *domain* of f and B the *codomain* of f . The **elements** of A are called inputs or *arguments* of f , and the **element** of B that is paired with an argument x by f is called the *value of f* for argument x , written $f(x)$.

The *range* $\text{ran}(f)$ of f is the subset of the codomain consisting of the values of f for some argument; $\text{ran}(f) = \{f(x) : x \in A\}$.

The diagram in **Figure fun.1** may help to think about functions. The ellipse on the left represents the function's *domain*; the ellipse on the right represents the function's *codomain*; and an arrow points from an *argument* in the domain to the corresponding *value* in the codomain.

Example fun.2. Multiplication takes pairs of natural numbers as inputs and maps them to natural numbers as outputs, so goes from $\mathbb{N} \times \mathbb{N}$ (the domain) to \mathbb{N} (the codomain). As it turns out, the range is also \mathbb{N} , since every $n \in \mathbb{N}$ is $n \times 1$.

Example fun.3. Multiplication is a function because it pairs each input—each pair of natural numbers—with a single output: $\times: \mathbb{N}^2 \rightarrow \mathbb{N}$. By contrast,




Figure fun.1: A function is a mapping of each **element** of one set to **an element** of another. An arrow points from an argument in the domain to the corresponding value in the codomain.

sfr:fun:bas:
fig:function

the square root operation applied to the domain \mathbb{N} is not functional, since each positive integer n has two square roots: \sqrt{n} and $-\sqrt{n}$. We can make it functional by only returning the positive square root: $\sqrt{} : \mathbb{N} \rightarrow \mathbb{R}$.

Example fun.4. The relation that pairs each student in a class with their final grade is a function—no student can get two different final grades in the same class. The relation that pairs each student in a class with their parents is not a function: students can have zero, or two, or more parents.

explanation

We can define functions by specifying in some precise way what the value of the function is for every possible argument. Different ways of doing this are by giving a formula, describing a method for computing the value, or listing the values for each argument. However functions are defined, we must make sure that for each argument we specify one, and only one, value.

Example fun.5. Let $f : \mathbb{N} \rightarrow \mathbb{N}$ be defined such that $f(x) = x + 1$. This is a definition that specifies f as a function which takes in natural numbers and outputs natural numbers. It tells us that, given a natural number x , f will output its successor $x + 1$. In this case, the codomain \mathbb{N} is not the range of f , since the natural number 0 is not the successor of any natural number. The range of f is the set of all positive integers, \mathbb{Z}^+ .

Example fun.6. Let $g : \mathbb{N} \rightarrow \mathbb{N}$ be defined such that $g(x) = x + 2 - 1$. This tells us that g is a function which takes in natural numbers and outputs natural numbers. Given a natural number n , g will output the predecessor of the successor of the successor of x , i.e., $x + 1$.

sfr:fun:bas:
examplefunext

explanation

We just considered two functions, f and g , with different *definitions*. However, these are the *same function*. After all, for any natural number n , we have that $f(n) = n + 1 = n + 2 - 1 = g(n)$. Otherwise put: our definitions for f and g specify the same mapping by means of different equations. Implicitly, then, we are relying upon a principle of extensionality for functions,

$$\text{if } \forall x f(x) = g(x), \text{ then } f = g$$

provided that f and g share the same domain and codomain.




Figure fun.2: A surjective function has every element of the codomain as a value.

sfr:fun:kin:
fig:surjective

Example fun.7. We can also define functions by cases. For instance, we could define $h: \mathbb{N} \rightarrow \mathbb{N}$ by

$$h(x) = \begin{cases} \frac{x}{2} & \text{if } x \text{ is even} \\ \frac{x+1}{2} & \text{if } x \text{ is odd.} \end{cases}$$

Since every natural number is either even or odd, the output of this function will always be a natural number. Just remember that if you define a function by cases, every possible input must fall into exactly one case. In some cases, this will require a proof that the cases are exhaustive and exclusive.

fun.2 Kinds of Functions

sfr:fun:kin:
sec It will be useful to introduce a kind of taxonomy for some of the kinds of explanation functions which we encounter most frequently.

To start, we might want to consider functions which have the property that every member of the codomain is a value of the function. Such functions are called surjective, and can be pictured as in Figure fun.2.

Definition fun.8 (Surjective function). A function $f: A \rightarrow B$ is surjective iff B is also the range of f , i.e., for every $y \in B$ there is at least one $x \in A$ such that $f(x) = y$, or in symbols:

$$(\forall y \in B)(\exists x \in A)f(x) = y.$$

We call such a function a surjection from A to B .

If you want to show that f is a surjection, then you need to show that every explanation object in f 's codomain is the value of $f(x)$ for some input x .

Note that any function induces a surjection. After all, given a function $f: A \rightarrow B$, let $f': A \rightarrow \text{ran}(f)$ be defined by $f'(x) = f(x)$. Since $\text{ran}(f)$ is defined as $\{f(x) \in B : x \in A\}$, this function f' is guaranteed to be a surjection

Now, any function maps each possible input to a unique output. But there explanation are also functions which never map different inputs to the same outputs. Such functions are called injective, and can be pictured as in Figure fun.3.



Figure fun.3: An injective function never maps two different arguments to the same value.

sfr:fun:kin:
fig:injective

Definition fun.9 (Injective function). A function $f: A \rightarrow B$ is *injective* iff for each $y \in B$ there is at most one $x \in A$ such that $f(x) = y$. We call such a function an *injection* from A to B .

explanation If you want to show that f is an *injection*, you need to show that for any elements x and y of f 's domain, if $f(x) = f(y)$, then $x = y$.

Example fun.10. The constant function $f: \mathbb{N} \rightarrow \mathbb{N}$ given by $f(x) = 1$ is neither *injective*, nor *surjective*.

The identity function $f: \mathbb{N} \rightarrow \mathbb{N}$ given by $f(x) = x$ is both *injective* and *surjective*.

The successor function $f: \mathbb{N} \rightarrow \mathbb{N}$ given by $f(x) = x + 1$ is *injective* but not *surjective*.

The function $f: \mathbb{N} \rightarrow \mathbb{N}$ defined by:

$$f(x) = \begin{cases} \frac{x}{2} & \text{if } x \text{ is even} \\ \frac{x+1}{2} & \text{if } x \text{ is odd.} \end{cases}$$

is *surjective*, but not *injective*.

explanation Often enough, we want to consider functions which are both *injective* and *surjective*. We call such functions *bijective*. They look like the function pictured in **Figure fun.4**. *Bijections* are also sometimes called *one-to-one correspondences*, since they uniquely pair elements of the codomain with elements of the domain.

Definition fun.11 (Bijection). A function $f: A \rightarrow B$ is *bijective* iff it is both *surjective* and *injective*. We call such a function a *bijection* from A to B (or between A and B).

fun.3 Functions as Relations

explanation A function which maps elements of A to elements of B obviously defines a relation between A and B , namely the relation which holds between x and y iff $f(x) = y$. In fact, we might even—if we are interested in reducing the building

sfr:fun:rel:
sec




Figure fun.4: A bijective function uniquely pairs the elements of the codomain with those of the domain.

sfr:fun:kin:
fig:bijective

blocks of mathematics for instance—*identify* the function f with this relation, i.e., with a set of pairs. This then raises the question: which relations define functions in this way?

Definition fun.12 (Graph of a function). Let $f: A \rightarrow B$ be a function. The *graph* of f is the relation $R_f \subseteq A \times B$ defined by

$$R_f = \{\langle x, y \rangle : f(x) = y\}.$$

The graph of a function is uniquely determined, by extensionality. Moreover, extensionality (on sets) will immediately vindicate the implicit principle of extensionality for functions, whereby if f and g share a domain and codomain then they are identical if they agree on all values.

explanation

Similarly, if a relation is “functional”, then it is the graph of a function.

sfr:fun:rel:
prop:graph-function

Proposition fun.13. Let $R \subseteq A \times B$ be such that:

1. If Rxy and Rxz then $y = z$; and
2. for every $x \in A$ there is some $y \in B$ such that $\langle x, y \rangle \in R$.

Then R is the graph of the function $f: A \rightarrow B$ defined by $f(x) = y$ iff Rxy .

Proof. Suppose there is a y such that Rxy . If there were another $z \neq y$ such that Rxz , the condition on R would be violated. Hence, if there is a y such that Rxy , this y is unique, and so f is well-defined. Obviously, $R_f = R$. \square

Every function $f: A \rightarrow B$ has a graph, i.e., a relation on $A \times B$ defined by $f(x) = y$. On the other hand, every relation $R \subseteq A \times B$ with the properties given in **Proposition fun.13** is the graph of a function $f: A \rightarrow B$. Because of this close connection between functions and their graphs, we can think of a function simply as its graph. In other words, functions can be identified with certain relations, i.e., with certain sets of tuples. We can now consider performing similar operations on functions as we performed on relations (see ??). In particular:

explanation

Definition fun.14. Let $f: A \rightarrow B$ be a function with $C \subseteq A$.

The *restriction* of f to C is the function $f|_C: C \rightarrow B$ defined by $(f|_C)(x) = f(x)$ for all $x \in C$. In other words, $f|_C = \{(x, y) \in R_f : x \in C\}$.

The *application* of f to C is $f[C] = \{f(x) : x \in C\}$. We also call this the *image* of C under f .

sfr:fun:rel:
defn:funimage

explanation It follows from these definition that $\text{ran}(f) = f[\text{dom}(f)]$, for any function f .

fun.4 Inverses of Functions

explanation We think of functions as maps. An obvious question to ask about functions, then, is whether the mapping can be “reversed.” For instance, the successor function $f(x) = x+1$ can be reversed, in the sense that the function $g(y) = y-1$ “undoes” what f does.

sfr:fun:inv:
sec

But we must be careful. Although the definition of g defines a function $\mathbb{Z} \rightarrow \mathbb{Z}$, it does not define a *function* $\mathbb{N} \rightarrow \mathbb{N}$, since $g(0) \notin \mathbb{N}$. So even in simple cases, it is not quite obvious whether a function can be reversed; it may depend on the domain and codomain.

This is made more precise by the notion of an inverse of a function.

Definition fun.15. A function $g: B \rightarrow A$ is an *inverse* of a function $f: A \rightarrow B$ if $f(g(y)) = y$ and $g(f(x)) = x$ for all $x \in A$ and $y \in B$.

If f has an inverse g , we often write f^{-1} instead of g .

explanation Now we will determine when functions have inverses. A good candidate for an inverse of $f: A \rightarrow B$ is $g: B \rightarrow A$ “defined by”

$$g(y) = \text{“the” } x \text{ such that } f(x) = y.$$

But the scare quotes around “defined by” (and “the”) suggest that this is not a definition. At least, it will not always work, with complete generality. For, in order for this definition to specify a function, there has to be one and only one x such that $f(x) = y$ —the output of g has to be uniquely specified. Moreover, it has to be specified for every $y \in B$. If there are x_1 and $x_2 \in A$ with $x_1 \neq x_2$ but $f(x_1) = f(x_2)$, then $g(y)$ would not be uniquely specified for $y = f(x_1) = f(x_2)$. And if there is no x at all such that $f(x) = y$, then $g(y)$ is not specified at all. In other words, for g to be defined, f must be both *injective* and *surjective*.

Let’s go slowly. We’ll divide the question into two: Given a function $f: A \rightarrow B$, when is there a function $g: B \rightarrow A$ so that $g(f(x)) = x$? Such a g “undoes” what f does, and is called a *left inverse* of f . Secondly, when is there a function $h: B \rightarrow A$ so that $f(h(y)) = y$? Such an h is called a *right inverse* of f —“undoes” what h does.

Proposition fun.16. If $f: A \rightarrow B$ is *injective*, then there is a left inverse $g: B \rightarrow A$ of f so that $g(f(x)) = x$ for all $x \in A$.

Proof. Suppose that $f: A \rightarrow B$ is **injective**. Consider a $y \in B$. If $y \in \text{ran}(f)$, there is an $x \in A$ so that $f(x) = y$. Because f is **injective**, there is only one such $x \in A$. Then we can define: $g(y) = x$, i.e., $g(y)$ is “the” $x \in A$ such that $f(x) = y$. If $y \notin \text{ran}(f)$, we can map it to any $a \in A$. So, we can pick an $a \in A$ and define $g: B \rightarrow A$ by:

$$g(y) = \begin{cases} x & \text{if } f(x) = y \\ a & \text{if } y \notin \text{ran}(f). \end{cases}$$

It is defined for all $y \in B$, since for each such $y \in \text{ran}(f)$ there is exactly one $x \in A$ such that $f(x) = y$. By definition, if $y = f(x)$, then $g(y) = x$, i.e., $g(f(x)) = x$. \square

Problem fun.1. Show that if $f: A \rightarrow B$ has a left inverse g , then f is **injective**.

Proposition fun.17. If $f: A \rightarrow B$ is **surjective**, then there is a right inverse $h: B \rightarrow A$ of f so that $f(h(y)) = y$ for all $y \in B$.

Proof. Suppose that $f: A \rightarrow B$ is **surjective**. Consider a $y \in B$. Since f is **surjective**, there is an $x_y \in A$ with $f(x_y) = y$. Then we can define: $h(y) = x_y$, i.e., for each $y \in B$ we choose some $x \in A$ so that $f(x) = y$; since f is **surjective** there is always at least one to choose from.¹ By definition, if $x = h(y)$, then $f(x) = y$, i.e., for any $y \in B$, $f(h(y)) = y$. \square

Problem fun.2. Show that if $f: A \rightarrow B$ has a right inverse h , then f is **surjective**.

By combining the ideas in the previous proof, we now get that every **bijection** has an inverse, i.e., there is a single function which is both a left and right inverse of f . explanation

[sfr:fun:inv](#)
[prop:bijection-inverse](#) **Proposition fun.18.** If $f: A \rightarrow B$ is **bijective**, there is a function $f^{-1}: B \rightarrow A$ so that for all $x \in A$, $f^{-1}(f(x)) = x$ and for all $y \in B$, $f(f^{-1}(y)) = y$.

Proof. Exercise. \square

Problem fun.3. Prove **Proposition fun.18**. You have to define f^{-1} , show that it is a function, and show that it is an inverse of f , i.e., $f^{-1}(f(x)) = x$ and $f(f^{-1}(y)) = y$ for all $x \in A$ and $y \in B$.

¹Since f is **surjective**, for every $y \in B$ the set $\{x : f(x) = y\}$ is nonempty. Our definition of h requires that we choose a single x from each of these sets. That this is always possible is actually not obvious—the possibility of making these choices is simply assumed as an axiom. In other words, this proposition assumes the so-called Axiom of Choice, an issue we will gloss over. However, in many specific cases, e.g., when $A = \mathbb{N}$ or is finite, or when f is **bijective**, the Axiom of Choice is not required. (In the particular case when f is **bijective**, for each $y \in B$ the set $\{x : f(x) = y\}$ has exactly one **element**, so that there is no choice to make.)




Figure fun.5: The composition $g \circ f$ of two functions f and g .

sfr:fun:cmp:
fig:composition

explanation There is a slightly more general way to extract inverses. We saw in **section fun.2** that every function f induces a **surjection** $f': A \rightarrow \text{ran}(f)$ by letting $f'(x) = f(x)$ for all $x \in A$. Clearly, if f is **injective**, then f' is **bijective**, so that it has a unique inverse by **Proposition fun.18**. By a very minor abuse of notation, we sometimes call the inverse of f' simply “the inverse of f .”

Proposition fun.19. *Show that if $f: A \rightarrow B$ has a left inverse g and a right inverse h , then $h = g$.*

sfr:fun:inv:
prop:left-right

Proof. Exercise. □

Problem fun.4. Prove **Proposition fun.19**.

Proposition fun.20. *Every function f has at most one inverse.*

sfr:fun:inv:
prop:inverse-unique

Proof. Suppose g and h are both inverses of f . Then in particular g is a left inverse of f and h is a right inverse. By **Proposition fun.19**, $g = h$. □

fun.5 Composition of Functions

explanation We saw in **section fun.4** that the inverse f^{-1} of a **bijection** f is itself a function. Another operation on functions is composition: we can define a new function by composing two functions, f and g , i.e., by first applying f and then g . Of course, this is only possible if the ranges and domains match, i.e., the range of f must be a subset of the domain of g .

sfr:fun:cmp:
sec

A diagram might help to explain the idea of composition. In **Figure fun.5**, we depict two functions $f: A \rightarrow B$ and $g: B \rightarrow C$ and their composition $(g \circ f)$. The function $(g \circ f): A \rightarrow C$ pairs each **element** of A with **an element** of C . We specify which **element** of C **an element** of A is paired with as follows: given an input $x \in A$, first apply the function f to x , which will output some $f(x) = y \in B$, then apply the function g to y , which will output some $g(f(x)) = g(y) = z \in C$.

Definition fun.21 (Composition). Let $f: A \rightarrow B$ and $g: B \rightarrow C$ be functions. The *composition* of f with g is $g \circ f: A \rightarrow C$, where $(g \circ f)(x) = g(f(x))$.

Example fun.22. Consider the functions $f(x) = x + 1$, and $g(x) = 2x$. Since $(g \circ f)(x) = g(f(x))$, for each input x you must first take its successor, then multiply the result by two. So their composition is given by $(g \circ f)(x) = 2(x+1)$.

Problem fun.5. Show that if $f: A \rightarrow B$ and $g: B \rightarrow C$ are both **injective**, then $g \circ f: A \rightarrow C$ is **injective**.

Problem fun.6. Show that if $f: A \rightarrow B$ and $g: B \rightarrow C$ are both **surjective**, then $g \circ f: A \rightarrow C$ is **surjective**.

Problem fun.7. Suppose $f: A \rightarrow B$ and $g: B \rightarrow C$. Show that the graph of $g \circ f$ is $R_f \mid R_g$.

fun.6 Partial Functions

sfr:fun:par:
sec It is sometimes useful to relax the definition of function so that it is not required explanation that the output of the function is defined for all possible inputs. Such mappings are called *partial functions*.

Definition fun.23. A *partial function* $f: A \rightarrow B$ is a mapping which assigns to every **element** of A at most one **element** of B . If f assigns an element of B to $x \in A$, we say $f(x)$ is *defined*, and otherwise *undefined*. If $f(x)$ is defined, we write $f(x) \downarrow$, otherwise $f(x) \uparrow$. The *domain* of a partial function f is the subset of A where it is defined, i.e., $\text{dom}(f) = \{x \in A : f(x) \downarrow\}$.

Example fun.24. Every function $f: A \rightarrow B$ is also a partial function. Partial functions that are defined everywhere on A —i.e., what we so far have simply called a function—are also called *total* functions.

Example fun.25. The partial function $f: \mathbb{R} \rightarrow \mathbb{R}$ given by $f(x) = 1/x$ is undefined for $x = 0$, and defined everywhere else.

Problem fun.8. Given $f: A \rightarrow B$, define the partial function $g: B \rightarrow A$ by: for any $y \in B$, if there is a unique $x \in A$ such that $f(x) = y$, then $g(y) = x$; otherwise $g(y) \uparrow$. Show that if f is injective, then $g(f(x)) = x$ for all $x \in \text{dom}(f)$, and $f(g(y)) = y$ for all $y \in \text{ran}(f)$.

Definition fun.26 (Graph of a partial function). Let $f: A \rightarrow B$ be a partial function. The *graph* of f is the relation $R_f \subseteq A \times B$ defined by

$$R_f = \{\langle x, y \rangle : f(x) = y\}.$$

Proposition fun.27. *Suppose $R \subseteq A \times B$ has the property that whenever Rxy and Rxy' then $y = y'$. Then R is the graph of the partial function $f: X \rightarrow Y$ defined by: if there is a y such that Rxy , then $f(x) = y$, otherwise $f(x) \uparrow$. If R is also serial, i.e., for each $x \in X$ there is a $y \in Y$ such that Rxy , then f is total.*

Proof. Suppose there is a y such that Rxy . If there were another $y' \neq y$ such that Rxy' , the condition on R would be violated. Hence, if there is a y such that Rxy , that y is unique, and so f is well-defined. Obviously, $R_f = R$ and f is total if R is serial. \square

Photo Credits

Bibliography